



This chapter covers the following EIGRP troubleshooting topics:

- Troubleshooting EIGRP neighbor relationships
- Troubleshooting EIGRP route advertisement
- Troubleshooting EIGRP route installation
- Troubleshooting EIGRP route flap
- Troubleshooting EIGRP route summarization
- Troubleshooting EIGRP route redistribution
- Troubleshooting EIGRP dial backup
- EIGRP error messages

Troubleshooting EIGRP

This chapter discusses some of the common problems in EIGRP and how to resolve those problems. Debugs, configurations, and useful **show** commands are also given where necessary.

NOTE

Debugs can be CPU-intensive and can adversely affect your network. Therefore, debugs are not recommended on a production network unless being instructed by Cisco's Technical Assistance Center (TAC).

Sometimes, there might be multiple causes for the same problem. Therefore, if one scenario doesn't fix the network problem, always check into other scenarios.

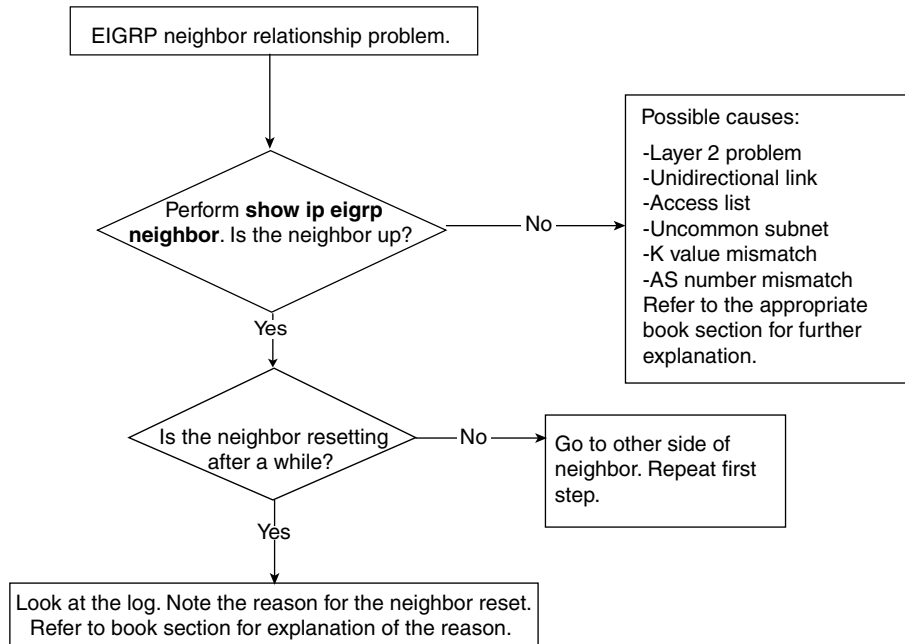
Troubleshooting EIGRP Neighbor Relationships

This section discusses methods of troubleshooting issues regarding EIGRP neighbor relationships. The following are the most common causes of problems with EIGRP neighbor relationships:

- Unidirectional link
- Uncommon subnet, primary, and secondary address mismatch
- Mismatched masks
- K value mismatches
- Mismatched AS numbers
- Stuck in active
- Layer 2 problem
- Access list denying multicast packets
- Manual change (summary router, metric change, route filter)

Figure 7-1 illustrates a general troubleshooting flowchart on EIGRP neighbor relationships.

Figure 7-1 *General Flowchart on Troubleshooting EIGRP Neighbor Relationships*



Consulting the EIGRP Log for Neighbor Changes

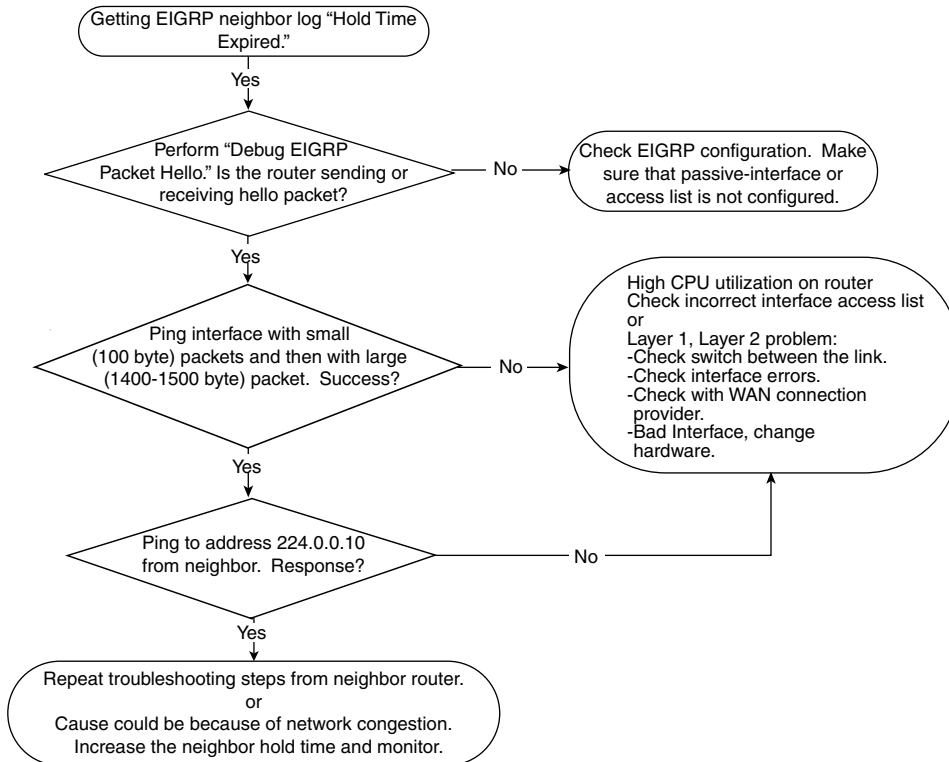
Whenever EIGRP resets its neighbor relationship, it is noted in the log with the reason for the reset. In the earlier Cisco IOS Software releases, configuration to enable this feature is required. The command **eigrp log-neighbor-change** is configured under router EIGRP. In Cisco IOS Software Release 12.1.3 and later, the **eigrp log-neighbor-change** command becomes the default setting for the router. An example of the EIGRP neighbor log looks something like this:

```
%DUAL-5-NBRCHANGE: IP-EIGRP EIGRP AS number: Neighbor neighbor IP address is down:
reason for neighbor down.
```

Table 7-1 documents the neighbor changes that you can find in the EIGRP log, along with the meaning and required action to fix the problem based on the log message.

Table 7-1 *Neighbor Changes Documented in the EIGRP Log*

Log Message	Meaning	Action for Troubleshooting
NEW ADJACENCY	Indicates that a new neighbor has been established.	No action is required.
PEER RESTARTED	Indicates that the other neighbor initiates the reset of the neighbor relationship. The router getting the message is not the one resetting the neighbor.	No action is required on the router that is getting the message. Gather EIGRP neighbor log information on the other neighbor.
HOLD TIME EXPIRED	Indicates that the router has not heard any EIGRP packets from the neighbor within the hold-time limit.	Because this is a packet-loss problem, check for a Layer 2 problem. Troubleshoot by using the flowchart shown in Figure 7-2.
RETRY LIMIT EXCEEDED	Indicates that EIGRP did not receive the acknowledgement from the neighbor for EIGRP reliable packets and that EIGRP already has tried to retransmit the reliable packet 16 times without any success.	Troubleshoot using the flowchart shown in Figure 7-3.
ROUTE FILTER CHANGED	Indicates that the EIGRP neighbor is resetting because there is a change in the route filter (distribute-list command under router EIGRP).	No action is needed. This is normal behavior in EIGRP, which needs to reset the neighbor when the route filter is changed, and to resynchronize the EIGRP topology table between neighbors.
INTERFACE DELAY CHANGED	Indicates that the EIGRP neighbor is resetting because there is a manual configuration change in the delay parameter on the interface.	No action is needed. This is normal behavior in EIGRP, which needs to reset the neighbor when the delay parameter is changed.
INTERFACE BANDWIDTH CHANGED	Indicates that the EIGRP neighbor is resetting because there is a manual configuration change in the interface bandwidth on the interface.	No action is needed. This is normal behavior in EIGRP, which needs to reset the neighbor when the bandwidth parameter is changed.
STUCK IN ACTIVE	Indicates that the EIGRP neighbor is resetting because EIGRP is stuck in active state. The neighbor getting reset is the result of stuck in active.	Troubleshoot from the stuck in active point of view. Refer to the section “EIGRP Neighbor Problem—Cause: Stuck in Active.”

Figure 7-2 Flowchart for Troubleshooting EIGRP Neighbor Relationship When Getting Neighbor Log Message HOLD TIME EXPIRED

EIGRP Neighbor Problem—Cause: Unidirectional Link

Sometimes, a problem with a WAN connection causes EIGRP to have a one-way neighbor relationship. A one-way neighbor relationship usually is caused by a unidirectional connection between the neighbors. The cause for unidirectional connection is usually a Layer 2 problem. For example, a link might be experiencing many CRC errors, a switch problem, or a **ping** test failure with large or small packets. In this case, you need a call to the group that is responsible for the link to check the integrity of the link. Sometimes, a simple misconfigured access list causes EIGRP to form a one-way neighbor relationship. Figure 7-4 illustrates an example of an EIGRP problem as a result of a unidirectional link.

Figure 7-3 Flowchart for Troubleshooting EIGRP Neighbor Relationship When Getting Neighbor Log **RETRY LIMIT EXCEEDED**

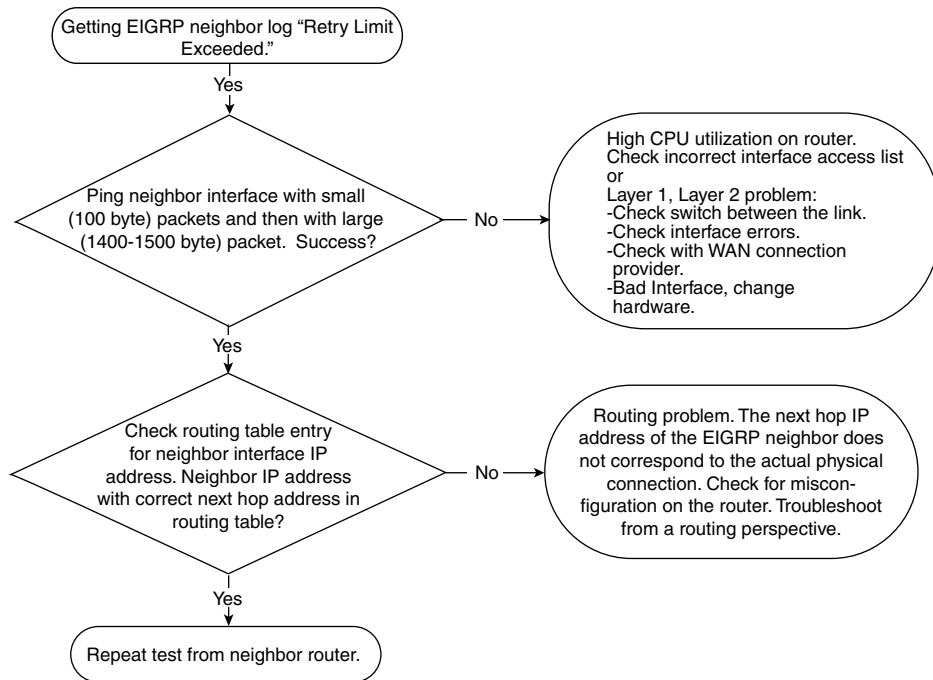



Figure 7-4 Network Topology Vulnerable to an EIGRP Neighbor Problem Because of a Unidirectional Link



In Figure 7-4, Routers RTR A and RTR B are connected by a WAN connection. The circuit from RTR A to RTR B is fine, but the circuit from RTR B to RTR A is broken. The results from the **show ip eigrp neighbor** command on RTR A will not show anything because RTR B's EIGRP hello packet can't make it to RTR A. Example 7-1 shows the output from **show ip eigrp neighbor** on RTR B.

Example 7-1 `show ip eigrp neighbors` Command Output on RTR B

RtrB#show ip eigrp neighbors							
IP-EIGRP neighbors for process 1							
H	Address	Interface	Hold	Uptime	SRTT (sec)	RTO (ms)	Q Cnt Seq Num
1	10.88.18.2	S0	14	00:00:15	0	5000	4 0

RTR B shows RTR A as a neighbor because RTR A's EIGRP hello packet has no problem reaching RTR B. From the output of the **show** command, the SRTT is at 0 ms, the retransmission timeout (RTO) timer is at 5000 ms, and the Q count is at 4 and is not decrementing. These three numbers give the biggest clue that this is a unidirectional link problem. The following is the meaning of SRTT, RTO, and Q count:

- **Smooth round-trip time (SRTT)**—The number of milliseconds it takes for an EIGRP packet to be sent to this neighbor and for the local router to receive an acknowledgment of that packet
- **Retransmission timeout (RTO), in milliseconds**—The amount of time that the software waits before retransmitting a packet from the retransmission queue to a neighbor
- **Q count**—The number of EIGRP packets (Update, Query, and Reply) that the software is waiting to send

Referring to Example 7-1, the fact that the SRTT timer is 0 indicates that no acknowledgement packets are being received. The Q count is not decrementing, which indicates that the router is trying to send EIGRP packets but no acknowledgement is being received. RTR B will retry 16 times to resend the packet; eventually, RTR B will reset the neighbor relationship with the log indicating **RETRY LIMIT EXCEEDED**, and the process starts again. Also, keep in mind that the 16 times retransmission of the same packet is done using unicast, not multicast. Therefore, the **RETRY LIMIT EXCEEDED** message indicates a problem with transmitting unicast packets over the link, and this is most likely a Layer 1 or Layer 2 problem.

The solution to this problem is to troubleshoot from a Layer 2 perspective. In this example, a call to the WAN provider is needed to find out why the circuit from RTR B to RTR A is broken. After the link between RTR B to RTR A is fixed, the problem will be resolved. Output from **show ip eigrp neighbors** in Example 7-2 shows that the neighbor relationship after the WAN link has been fixed.

Example 7-2 `show ip eigrp neighbors` Command Output Confirms Problem Resolution

RtrB#show ip eigrp neighbors							
IP-EIGRP neighbors for process 1							
H	Address	Interface	Hold	Uptime	SRTT (sec)	RTO (ms)	Q Cnt Seq Num
1	10.88.18.2	S0	14	01:26:30	149	894	0 291

Notice that the Q count column is 0 and that the SRTT and RTO have valid values now.

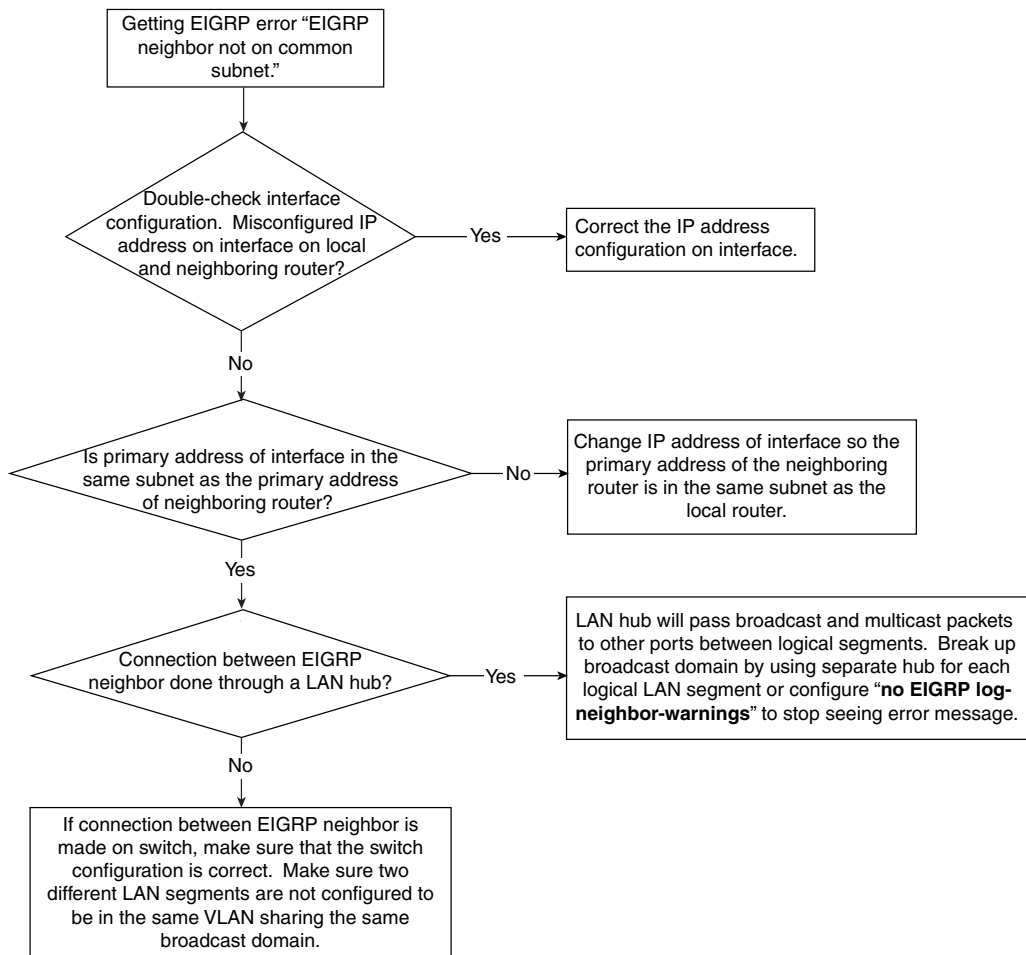
EIGRP Neighbor Problem—Cause: Uncommon Subnet

Many times, EIGRP won't establish neighbor relationships because the neighbors are not in the same subnet. Usually, the cause of this problem is router misconfiguration. When EIGRP has problems establishing neighbor relationships because of an uncommon subnet, the following error message appears:

IP-EIGRP: Neighbor *ip address* not on common subnet for *interface*

Figure 7-5 shows the flowchart for troubleshooting the problem when the "Neighbor not on common subnet" error appears on the router.

Figure 7-5 Problem-Resolution Flowchart



According to the troubleshooting flowchart in Figure 7-5, the three causes of getting the “EIGRP neighbor not on common subnet” error message are the following:

- The IP address has been misconfigured on interfaces.
- The primary and secondary IP addresses of the neighboring interface don’t match.
- A switch or hub between the EIGRP neighbor connection is misconfigured or is leaking multicast packet to other ports.


Misconfiguration of the IP Address on the Interfaces

Sometimes, an EIGRP neighbor that is not on a common subnet with other EIGRP neighbors is simply the result of misconfiguring the IP address on the interfaces. For example, the network administrator might mistype IP address 192.168.3.1 255.255.255.252 as 192.168.3.11 255.255.255.252, which causes EIGRP to complain about the neighbor not being on a common subnet.

Primary and Secondary IP Addresses of the Neighboring Interface Don’t Match

As mentioned in Chapter 6, “Understanding Enhanced Interior Gateway Routing Protocol (EIGRP),” EIGRP sources the hello packet from the primary address of the interface. If the primary network address on one router is used as a secondary network address on the second router, and vice versa, no neighbor relationship will be formed and the routers will complain about the neighbor not being on a common subnet. Figure 7-6 illustrates such a scenario.

Figure 7-6 *Network Topology Vulnerable to EIGRP Neighbor Problems Because of Primary and Secondary IP Address Mismatch*



In Figure 7-6, Router A and Router B have a primary address in the 10.1.1.0/24 network range, while Router C has an address range of 50.1.1.0/24 configured. When Router A or Router B sends out the EIGRP hello packet, the source of the hello packet will be either 10.1.1.1 or 10.1.1.2, depending on which router sends out the hello. When Router C receives the hello packet from Router A or Router B, it notices that the source is from the 10.1.1.0 network. Because Router C has an IP address of 50.1.1.3 configured on the interface, Router C will not process the hello packet from Router A or Router B because they are from a different network. Therefore, no neighbor relationship is formed from Router C to either Router A or Router B.

The solution for this example is to match all the IP addresses on the segment to the primary address space. For the network in Figure 7-6, you need to configure Router C to be in the primary address space of 10.1.1.0/24.


Switch or Hub Between EIGRP Neighbor Connection Is Misconfigured or Is Leaking Multicast Packets to Other Ports

If the IP address configuration is correct on the interface between EIGRP neighbors, you might want to check the configuration on the switch or the hub that connects the EIGRP neighbors. If a single LAN hub connects the EIGRP neighbors for different LAN segment, the hub passes broadcast and multicast packets to other ports between two logical LAN segments. So, the multicast EIGRP hello from LAN segment 1 will be seen on the neighbor located in LAN segment 2 if a single hub connects all the LAN devices on different LAN segments. The solution is to break up the broadcast domain by using a separate hub for each LAN segment or simply configuring **no eigrp log-neighbor-warnings** under EIGRP configuration to stop seeing the error message.

If a LAN switch connects the LAN devices, you might want to check the configuration of the switch. Make sure that the switch is not configured so that different LAN segments reside within the same VLAN. Make sure that the switch is configured so that each LAN segment has its own broadcast domain and does not share its broadcast domain with other LAN segments.

EIGRP Neighbor Problem—Cause: Mismatched Masks

Sometimes, a simple misconfiguration on the interface subnet mask causes an EIGRP neighbor problem. Figure 7-7 illustrates a network diagram for such a scenario.

Figure 7-7 Network Topology Vulnerable to EIGRP Neighbor Problems Because of Mismatched Masks

Example 7-3 shows the configuration for Routers A, B, and C.

Example 7-3 Router A, B, and C Configurations for the Network in Figure 7-7

```

Router A#interface serial 0
ip address 10.1.1.2 255.255.255.128
interface serial 1
ip address 10.1.3.1 255.255.255.0

Router B#interface serial 0
ip address 10.1.1.1 255.255.255.0
interface ethernet 0
ip address 10.1.2.1 255.255.255.0

Router C#interface ethernet 0
ip address 10.1.2.2 255.255.255.0
interface serial 0
ip address 10.1.3.2 255.255.255.0

```

Notice the mismatched mask on the serial interface of Router A and Router B. Router A has a mask of 255.255.255.128, while Router B has a mask of 255.255.255.0 on Serial 0. Initially, EIGRP has no problem forming the neighbor between Router A and Router B because 10.1.1.1 and 10.1.1.2 are in the same subnet. The problem occurs when a neighbor relationship is established and Router A and Router B begin to exchange EIGRP topology tables and install routes based on the EIGRP topology table, as demonstrated in Example 7-4.

Example 7-4 *Routing Tables from Router B and Router C*

```

Router B#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
        i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
        * - candidate default, U - per-user static route, o - ODR
        P - periodic downloaded static route

Gateway of last resort is not set
C 10.1.1.0/24   Serial 0
D 10.1.1.0/25   10.1.2.2

Router c#show ip route eigrp
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
        i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
        * - candidate default, U - per-user static route, o - ODR
        P - periodic downloaded static route

Gateway of last resort is not set
D 10.1.1.0/24   10.1.2.1
D 10.1.1.0/25   10.1.3.1

```

When Router B sends Router A an EIGRP update, Router A responds to the update with an EIGRP acknowledgement packet with a destination address of 10.1.1.1 to Router B. When Router B receives the packet, it forwards the ACK packet to Router C instead of processing it because Router B has a more specific route from Router C. Router B has a more specific route of 10.1.1.0/25 with the next hop to 10.1.2.2. This /25 route overrides the /24 route because /25 is more specific than /24. When Router C receives the ACK packet from Router B, it looks at its routing table for the 10.1.1.1 entry, and the routing table points to Router A. Router C then forwards the ACK packet back to Router A. This creates a routing loop. The packet to 10.1.1.1 loops from Router A to Router B, from Router B to Router C, and back from Router C to Router A. As a result, Router B won't process the ACK packet from Router A; Router B will think that Router A never ACK'ed the update packet, and Router B will reset the neighbor after 16 retries.


The solution for this problem: Configure the right subnet mask on Router A's Serial 0 interface to 255.255.255.0.

EIGRP Neighbor Problem—Cause: Mismatched K Values

For EIGRP to establish its neighbors, the K constant value to manipulate the EIGRP metric must be the same. Refer to Chapter 6 for an explanation of the K values. In EIGRP's

metric calculation, the default for the K value is set so that only the bandwidth and the delay of the interface are used to calculate the EIGRP metric. Many times, the network administrator might want other interface factors, such as load and reliability, to determine the EIGRP metric. Therefore, the K values are changed. Because only bandwidth and delay are used in calculations, the remaining K values are set to a value of 0 by default. However, the K values must be the same for all the routers, or EIGRP won't establish a neighbor relationship. Figure 7-8 shows an example of this case.

Figure 7-8 *Network Vulnerable to EIGRP Neighbor Problems Because of Mismatched K Values*



For the network in Figure 7-8, K1 is bandwidth and K3 is delay. The network administrator changed the K values of RTR B to all 1s from K1 to K4, while RTR A retains the default value of K1 and K3 to be 1. In this example, RTR A and RTR B will not form EIGRP neighbor relationship because the K values don't match. Example 7-5 shows the configuration for RTR B.

Example 7-5 *Configuration for RTR B in Figure 7-8*

```
RTR B#router eigrp 1
network xxxx
metric weights 0 1 1 1 1 0
```

RTR B's configuration includes the extra **metric weights** command. The first number is the type of service (ToS) number, which, because it's not supported, gets a value of 0. The five numbers after the ToS are the K1 through K5 values.

Troubleshooting this problem requires careful scrutiny of the router's configuration. The solution for this problem is to change all the K values to be the same on all the neighboring routers. In this example, in Router A, changing the K values to match the K value of Router B will solve the problem, as demonstrated in Example 7-6.


Example 7-6 *Configuring the K Values on Router A to Match Router B*

```
RTR A#router eigrp 1
network xxxx
metric weights 0 1 1 1 1 0
```

EIGRP Neighbor Problem—Cause: Mismatched AS Number

EIGRP won't form any neighbor relationships with neighbors in different autonomous systems. If the AS numbers are mismatched, no adjacency is formed. This problem is usually caused by misconfiguration on the routers. Figure 7-9 illustrates such a problem.

Figure 7-9 Network Experiencing an EIGRP Neighbor Problem Because Mismatched AS Numbers



In the network shown in Figure 7-9, RTR A and RTR B are in the EIGRP AS number of 1 and the proper network numbers have been configured; however, no EIGRP neighbor relationship is formed between RTR A and RTR B. Begin by checking the configuration of RTR A and RTR B in Example 7-7.

Example 7-7 Configurations for RTR A and RTR B in Figure 7-9

```
RTR B#show running-config
interface serial 0
IP address 10.1.1.1 255.255.255.0
router eigrp 11
network 10.0.0.0
```

```
RTR A#show running-config
Interface serial 0
IP address 10.1.1.2 255.255.255.0
router eigrp 1
network 10.0.0.0
```

You should notice the misconfiguration immediately. RTR B's Serial 0 interface is configured to be in EIGRP AS number 11, while RTR A's Serial 0 is configured to be in EIGRP AS number 1. Because the AS numbers don't match across the link, no EIGRP neighbor relationship will be formed. To resolve this problem, simply configure both routers with the same EIGRP AS number, as shown in Example 7-8. In this example, both routers will be configured to be in EIGRP AS 1.

Example 7-8 Configuring Both Routers with the Same EIGRP AS Numbers

```
RTR A#router eigrp 1
network 10.0.0.0
```

```
RTR B#router eigrp 1
network 10.0.0.0
```

EIGRP Neighbor Problem—Cause: Stuck in Active

Sometimes, EIGRP resets the neighbor relationship because of a “stuck in active” condition. The error message is

```
%DUAL-3-SIA: Route network mask stuck-in-active state in IP-EIGRP AS. Cleaning up
```

This section discusses the method of troubleshooting the EIGRP stuck in active error.

Reviewing the EIGRP DUAL Process


To resolve an EIGRP stuck in active error, you need to understand the DUAL process in EIGRP. Refer to Chapter 6 for thorough coverage of the DUAL process, although it is reviewed here as well.

EIGRP is an advanced distance-vector protocol; it doesn't have LSA flooding, like OSPF, or a link-state protocol to tell the protocol the overall view of the network. EIGRP relies only on its neighbors for information on network reachability and availability. EIGRP keeps a list of backup routes called *feasible successors*. When the primary route is not available, EIGRP immediately uses the feasible successor as the backup route. This shortens convergence time. Now, if the primary route is gone and no feasible successor is available, the route is in active state. The only way for EIGRP to converge quickly is to query its neighbors about the unavailable route. If the neighbor doesn't know the status of the route, the neighbor asks its neighbors, and so on, until the edge of the network is reached. The query stops if one of the following occurs:

- All queries are answered from all the neighbors.
- The end of network is reached.
- The lost route is unknown to the neighbors.

The problem is that, if there are no query boundaries, EIGRP potentially can ask every router in the network for a lost route. When EIGRP first queries its neighbor, a stuck in active timer starts. By default, the timer is three minutes. If, in three minutes, EIGRP doesn't receive the query response from all its neighbors, EIGRP declares that the route is stuck in active state and resets the neighbor that has not responded to the query. Figure 7-10 illustrates the query process of EIGRP when a route is lost.

Figure 7-10 *Illustration of EIGRP Query Process When a Route Is Lost*



In Figure 7-10, Router A lost its Ethernet interface. Because it doesn't have a feasible successor, the route becomes active and Router A queries its neighbors, Router B and Router C. Now, Router B doesn't know how to reach the lost network, so it asks its neighbors, Router D and Router E. Similarly, Router C asks its neighbors, Router F and Router G. Because Routers D, E, F, and G also don't know how to reach the lost network, they query the downstream neighbors. At this point, the edge of the network is reached and the edge router doesn't have any more neighbors to query. The edge router then replies back to Routers D, E, F, and G. Those routers reply back to Routers B and C, and finally to Router A. The query process then stops. Figure 7-10 shows the cascade effect of the EIGRP query process, in which the query travels from the original router to the edge of the network and back to the original router.

Determining Active/Stuck in Active Routes with **show ip eigrp topology active**

You must answer two questions to troubleshoot the EIGRP stuck in active problem:

- Why is the route active?
- Why is the route stuck?

Determining why the route is active is not a difficult task. Sometimes, the route that constantly is going active could be due to flapping link. Or, if the route is a host route (/32 route), it's possible that it is from a dial-in connection that gets disconnected. However, trying to determine why the active route becomes stuck is a much harder task—and more important to learn. Usually, an active route gets stuck for one of the following reasons:

- Bad or congested links
- Low router resources, such as low memory or high CPU on the router
- Long query range
- Excessive redundancy

By default, the stuck in active timer is only three minutes. In other words, if the EIGRP neighbor doesn't hear a reply for the query in three minutes, neighbors are reset. This adds difficulty in troubleshooting EIGRP stuck in active because every time an active route is stuck, you have only three minutes to track down the active route query path and hopefully find the cause.

The tool that you need to troubleshoot the EIGRP stuck in active error is the **show ip eigrp topology active** command. This command shows what routes are currently active, how long the routes have been active, and which neighbors have and have not replied to the query. From the output, you can determine which neighbors have not replied to the query, and you can track the query path and find out the status of the query by hopping to the neighbors that have not replied. Example 7-9 shows sample output from the **show ip eigrp topology active** command.

Example 7-9 *Sample Output of show ip eigrp topology active Command*

```
Router#show ip eigrp topology active
IP-EIGRP Topology Table for AS(1)/ID(10.1.4.2)
A 20.2.1.0/24, 1 successors, FD is Inaccessible, Q
  1 replies, active 00:01:43, query-origin: Successor Origin
    via 10.1.3.1 (Infinity/Infinity), Serial1/0
    via 10.1.4.1 (Infinity/Infinity), Serial1/1, serno 146
  Remaining replies:
    Via 10.1.5.2, R, Serial1/2
```

As the output in Example 7-9 indicates, the route for 20.2.1.0 is in active state and has been active for 1 minute and 43 seconds. query-origin is Successor Origin, which means that this route's successor sends the query to this router. At this point, it has gotten replies from 10.1.3.1 and 10.1.4.1; the reply is infinity, which means that these two routers also don't know about the route 20.2.1.0. The most important output of the **show ip eigrp topology**

active command is the Remaining replies: section. From the output of Example 7-9, this router shows that the neighbor 10.1.5.2 from interface Serial1/2 has not replied to the query.

To proceed further with troubleshooting, you must Telnet to the 10.1.5.2 router to see the status of its EIGRP active routes using the same command, **show ip eigrp topology active**. Sometimes, the router does not list the neighbors that have not replied to the queries under the Remaining replies: section. Example 7-10 shows another output of **show ip eigrp topology active**.

Example 7-10 Another Sample Output of the **show ip eigrp topology active** Command

```
Router#show ip eigrp topology active
IP-EIGRP Topology Table for AS(110)/ID(175.62.8.1)
A 11.11.11.0/24, 1 successors, FD is Inaccessible
  1 replies, active 00:02:06, query-origin: Successor Origin
    via 1.1.1.2 (Infinity/Infinity), r, Serial1/0, serno 171
    via 10.1.1.2 (Infinity/Infinity), Serial1/1, serno 173
```

In Example 7-10, the only difference in output from Example 7-9 is the list of neighbors that have not replied to the router. However, this doesn't mean that all of the neighbors have replied to the queries. In Example 7-10, neighbor 1.1.1.2 has an **r** next to the address of 1.1.1.2. This also means that the neighbor has not replied to the queries. In other words, the router has two ways of representing neighbors that have not replied to the queries. One is to have them listed under the Remaining replies: section; the other is to have an **r** next to the neighbor interface IP address. When using the **show ip eigrp topology active** command, the router can use any combination of these methods to represent neighbors that have not yet replied to the queries, as demonstrated in Example 7-11.

Example 7-11 Output of **show ip eigrp topology active** That Shows a Combination Representation of Neighbors That Have Not Replied to the Queries

```
Router#show ip eigrp topology active
IP-EIGRP Topology Table for AS(110)/ID(175.62.8.1)
A 11.11.11.0/24, 1 successors, FD is Inaccessible
  1 replies, active 00:02:06, query-origin: Successor Origin
    via 1.1.1.2(Infinity/Infinity),r, Serial1/0, serno 171
    via 10.1.1.2 (Infinity/Infinity), Serial1/1, serno 173
  Remaining replies:
    via 10.1.5.2, r, Serial1/2
```

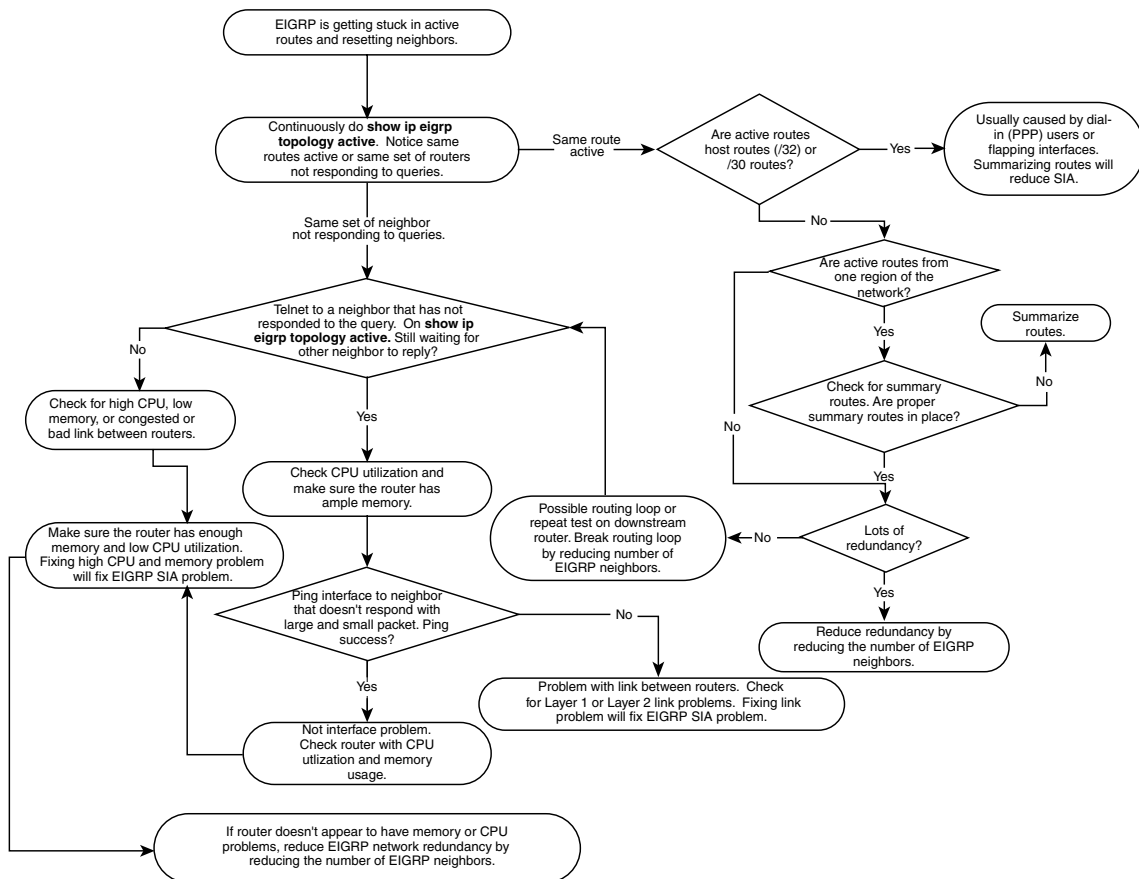
In Example 7-11, the neighbors that have not replied to the queries are 1.1.1.2 and 10.1.5.2. Only one of the nonreplying neighbors 10.1.5.2 is listed under the Remaining replies: section; the other neighbor, 1.1.1.2, that has not replied is listed with the other replying neighbor. To summarize, when issuing the **show ip eigrp topology active** command, the most important part to look for is the neighbors that have not replied to the query. To look for such a neighbor, look for neighbors that have the **r** next to their interface IP addresses.

Methodology for Troubleshooting the Stuck in Active Problem

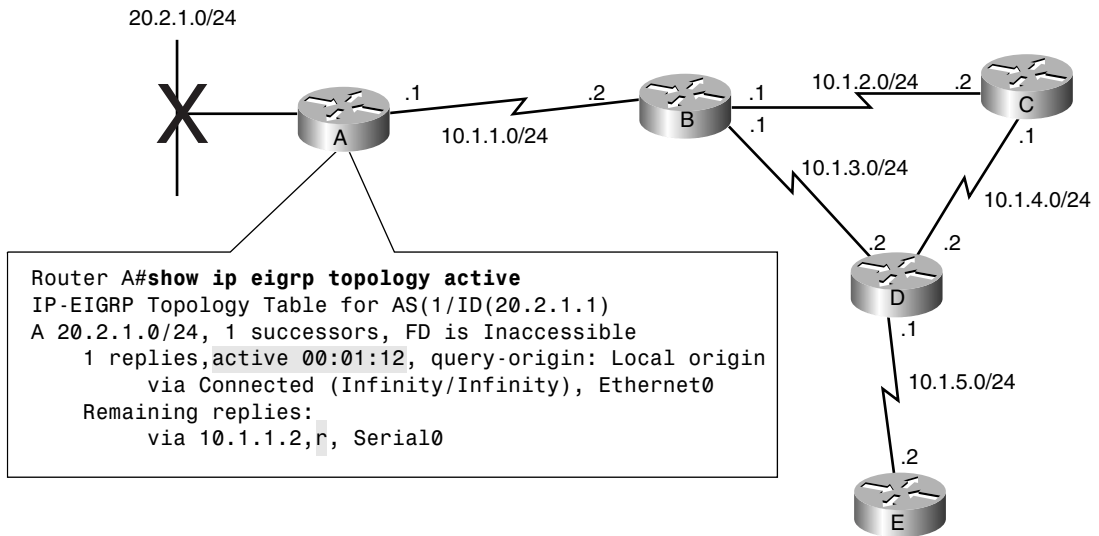
The methods for troubleshooting an EIGRP stuck in active problem and the **show ip eigrp topology active** command are useful only when the problem is happening. When the stuck in active event is over and the network stabilizes, it is extremely difficult, if not impossible, to backtrack the problem and find out the cause.

Figure 7-11 shows the flowchart for troubleshooting the EIGRP stuck in active problem.

Figure 7-11 Flowchart for Resolving the EIGRP Stuck in Active Problem



Consider the network shown in Figure 7-12 for an example of troubleshooting the EIGRP stuck in active problem.

Figure 7-12 Network Topology for EIGRP Stuck in Active Troubleshooting Example

In Figure 7-12, Router A has an Ethernet interface with network 20.2.1.0/24 that just went away. Router A doesn't have a feasible successor to go to as a backup route. Router A has no choice but to put the 20.2.1.0/24 route into active state and query its neighbor, Router B. Notice the output of **show ip eigrp topology active** in Router A. The 20.2.1.0/24 route has gone active for 1 minute and 12 seconds, and the neighbor that has not responded is listed as 10.1.1.2 from Serial0, which is Router B. The next step is to Telnet to Router B to see the active route status in Router B. Figure 7-13 shows the active route status in Router B by performing the command **show ip eigrp topology active**.

In Figure 7-13, the command **show ip eigrp topology active** on Router B shows that the route 20.2.1.0/24 is also in active status in Router B and that it has gone active for 1 minute and 23 seconds. Most importantly, Router B can't reply to Router A about route 20.2.1.0/24 because Router B is still waiting for the neighbor with IP address of 10.1.3.2 (Router D) from Serial1/2 to reply to the query. The next step is to go to Router D to see the status of the active route 20.2.1.0/24 and see why Router D has not replied to the query. Figure 7-14 shows the output of **show ip eigrp topology active** on Router D.

Router D also put the route 20.2.1.0/24 in active state, and it has been in active state for 1 minute and 43 seconds. Router D can't answer Router B's query because Router D is waiting for the router with the IP address of 10.1.5.2 from Serial1/2 (Router E) to respond to the query. The next step is to go to Router E to see the status of the active route 20.2.1.0/24 and to find out why Router E is not replying to the query. Figure 7-15 shows the status of the active route on Router E.

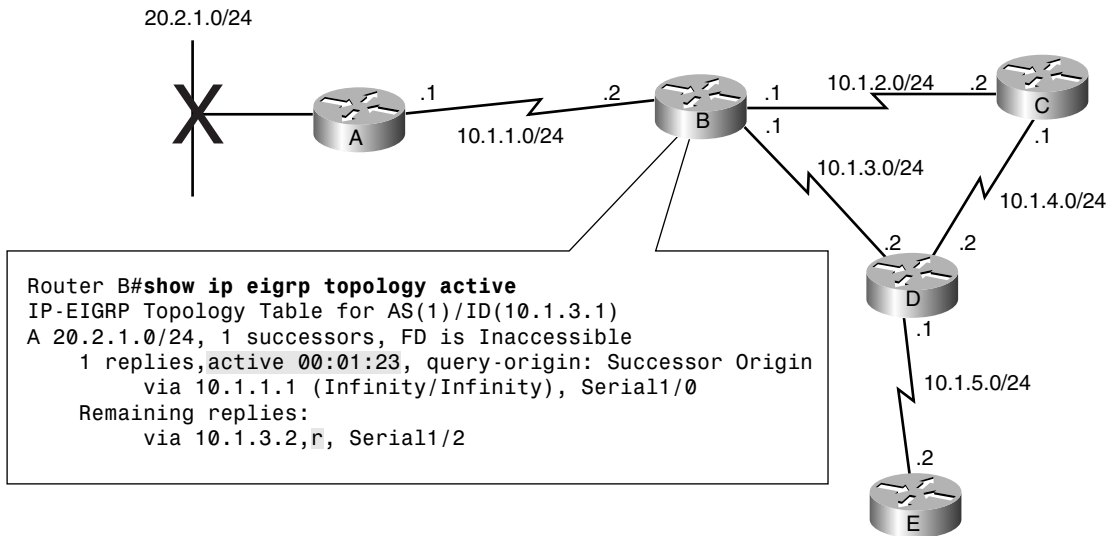
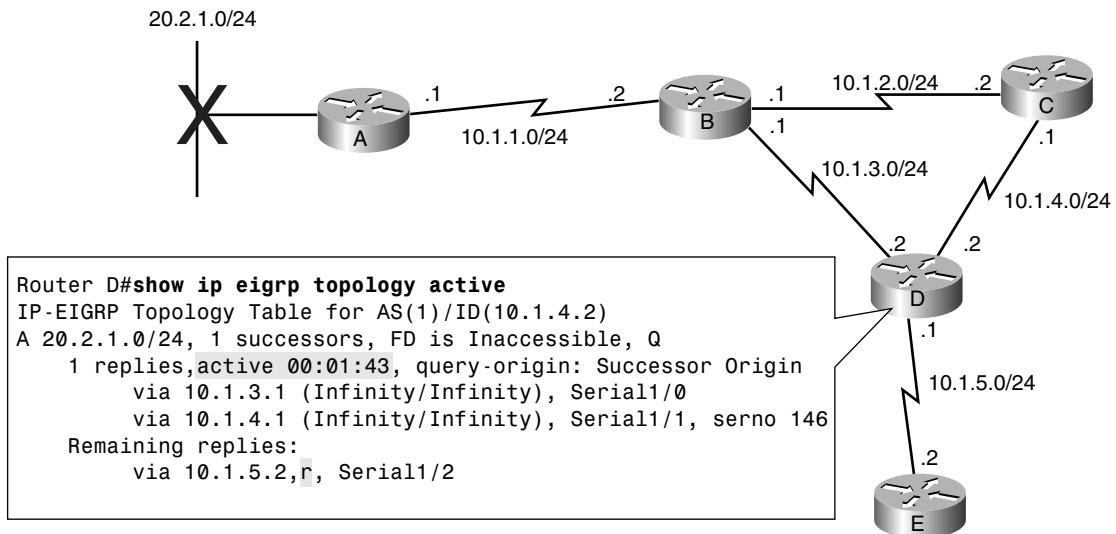

Figure 7-13 Active Route Status on Router B for Troubleshooting EIGRP Stuck in Active Example**Figure 7-14** Active Route Status on Router D for Troubleshooting EIGRP Stuck in Active Example

Figure 7-15 Active Route Status on Router E for the Troubleshooting EIGRP Stuck in Active Example



The output for **show ip eigrp topology active** didn't show anything for Router E. This indicates that, as far as Router E is concerned, there are no routes in active state. Now you should Telnet back to Router D to double-check whether the router is still in the active state for route 20.2.1.0/24. Telnetting back to Router D shows that Router D is still in active state for route 20.2.1.0/24, but Router E doesn't have any routes in active state. What's going on?

To summarize what has been going on so far, the chain of event is as follows:

- 1 Router A went active for route 20.2.1.0/24 and is waiting for Router B to reply to the query.
- 2 Router B can't reply because it is waiting for Router D's query response.
- 3 Router D can't reply because it is waiting for Router E to reply to the query.
- 4 Finally, the **show ip eigrp topology active** command in Router E shows that Router E does not think that any routes are active, while going back to Router D shows that the route 20.2.1.0/24 is still in active state.

From this sequence of events, you can see that there is clearly a discrepancy between Router D and Router E. More investigation is needed between these routers.

A look at Router D and Router E's router CPU utilization and memory usage doesn't show a problem. Both routers' CPU utilization and available memory are normal. You need to look at Router D's neighbor list to see if there is a problem with the neighbors. Example 7-12 shows Router D's EIGRP neighbor list.

Example 7-12 Router D's EIGRP Neighbor List

RTRD#show ip eigrp neighbors								
IP-EIGRP neighbors for process 1								
H	Address	Interface	Hold	Uptime (sec)	SRTT (ms)	RT0	Q Cnt	Seq Num
2	10.1.5.2	Se1/2	13	00:00:14	0	5000	1	0
1	10.1.3.1	Se1/0	13	01:22:54	227	1362	0	385
0	10.1.4.1	Se1/1	10	01:24:08	182	1140	0	171

From Example 7-12, notice that there is a problem in Router D with EIGRP sending a reliable packet to the neighbor with IP address of 10.1.5.2 (Router E). The Q count is 1, and performing the **show ip eigrp neighbors** command a few times in succession shows that the Q count is not decrementing.

The RTO counter is at its maximum value of 5000 ms. This indicates that Router D is trying to send a reliable packet to Router E, but Router E never acknowledges the reliable packet back to Router D. Because Router E doesn't appear to have a high CPU or memory problem, you should test the link reliability between Router D and Router E. Now send five **ping** packets from Router D to IP address 10.1.5.2 (Router E's serial interface) to see what happens. Example 7-13 shows the result of the **ping** test.

Example 7-13 Result of **ping** Test from Router D to Router E

Router D#ping 10.1.5.2	
Type escape sequence to abort.	
Sending 5, 100-byte ICMP Echos to 10.1.5.2, timeout is 2 seconds:	
.....	
Success rate is 0 percent (0/5)	

The **ping** test in Example 7-13 shows the success rate is 0 percent. This test shows that a link problem exists between Router D and Router E. The link is capable of passing a multicast packet to establish an EIGRP neighbor relationship, but it is having problems transmitting a unicast packet. This link problem is the root cause of the EIGRP stuck in active problem in this example. The way to troubleshoot the EIGRP stuck in active problem is to chase hop by hop the query path and find out the status of active route at each hop.

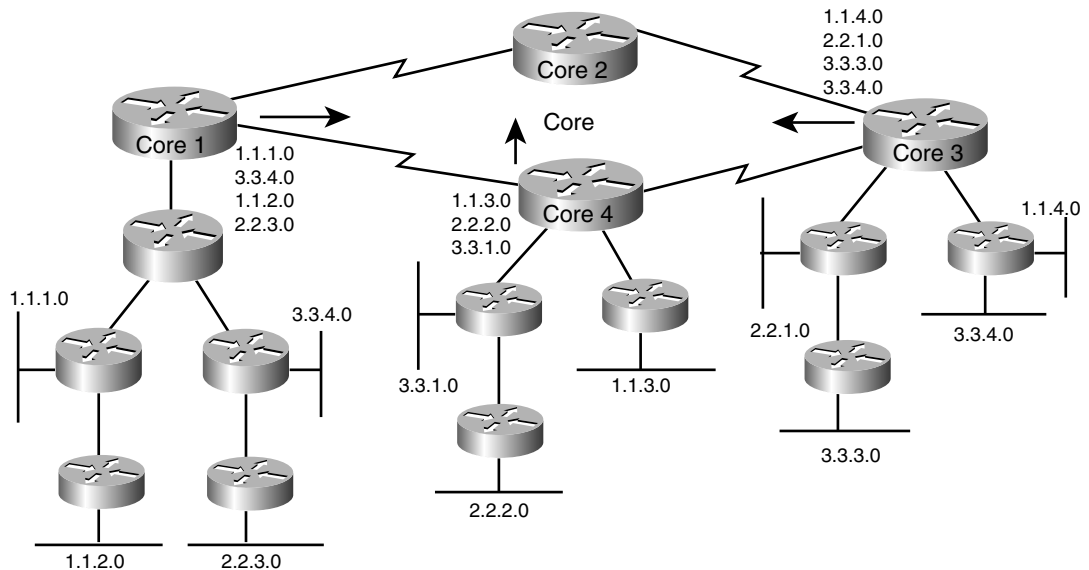
The aforementioned process is typical troubleshooting methodology for combatting the EIGRP stuck in active problem.

Sometimes, chasing the query path hop by hop leads to a loop, or there are simply too many neighbors that didn't reply to the query. In this case, simplify and reduce the complexity of the EIGRP topology by cutting down the redundancy. The simpler the EIGRP topology is, the simpler it is to troubleshoot an EIGRP stuck in active problem.


The ultimate solution for preventing the EIGRP stuck in active problem is to manually summarize the routes whenever possible and to have a hierarchical network design. The more network EIGRP summarizes, the less work EIGRP has to do when a major convergence takes

place. Therefore, this reduces the number of queries being sent out and ultimately reduces the occurrence of an EIGRP stuck in active error. Figure 7-16 shows an example of a poor network design that will not scale in a large EIGRP network.

Figure 7-16 *Example of a Nonscalable EIGRP Network*



In Figure 7-16, each core router represents a region of the entire network and shows that there is no hierarchy in IP addressing scheme. The Core 1 router is injecting routes 1.1.1.0, 3.3.4.0, 1.1.2.0, and 2.2.3.0 into the core network. The addresses are so scattered that no manual summarization is possible. The other core routers are experiencing the same problem. The Core 3 and Core 4 routers can't summarize any routes into the core network. As a result, if the Ethernet link of the 3.3.3.0 network keeps flapping, the query would travel to the Core 3 router and then the query also would be seen in the Core 1 and Core 4 region. Ultimately, the query will traverse to all the routers in the internetwork; this would dramatically increase the likelihood of an EIGRP stuck in active problem. The best practice is to readdress the IP address scheme. One region should take only a block of IP addresses; this way, the core routers would be capable of summarizing the routes into the core, resulting in a reduced routing table in the core. The routers and the query would be contained only in one region. Figure 7-17 shows an improved and more scalable EIGRP network design.

Figure 7-17 Scalable EIGRP Network Design Improvement on Network in Figure 7-16

Comparing Figures 7-16 and 7-17, you can see that the network presented in Figure 7-17 is more structured. The Core 1 router region takes only the 1.0.0.0 block of IP addresses, the Core region 4 takes only the 2.0.0.0 block, and Core 3 region takes only the 3.0.0.0 block of IP addresses. This enables the three core routers to summarize their routes into the core. If the Ethernet network of 3.3.3.0 flaps in the Core 3 region, the query would be bounded only in the Core 3 region and would not travel the entire network to affect all the routers in the network. Summarization and hierarchy are the best design practices for a large-scale EIGRP network.

Troubleshooting EIGRP Route Advertisement

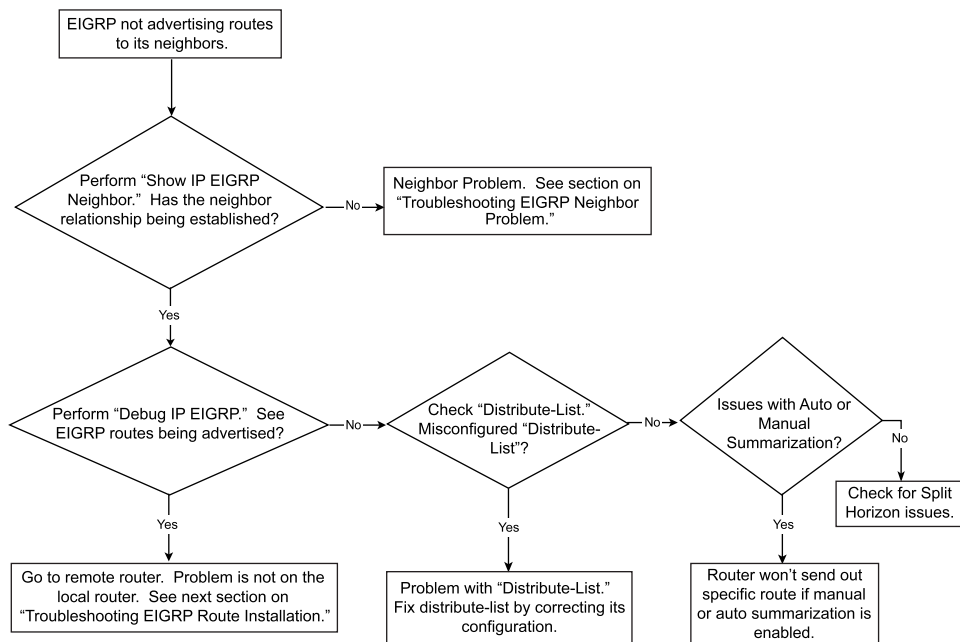
Sometimes, EIGRP has issues with route advertisement. This section discusses methods for troubleshooting EIGRP route advertisement problems, which can be categorized as follows:

- EIGRP is not advertising routes to neighbors when the network administrators think that it should.
- EIGRP is advertising routes to neighbors when the network administrators think that it shouldn't.
- EIGRP is advertising routes with a metric that is not understood by the network administrators.

EIGRP Is Not Advertising Routes to Neighbors When the Network Administrators Think That It Should

This section discusses methods for troubleshooting issues related to EIGRP not advertising routes to the neighbors. Figure 7-18 shows a flowchart documenting how to troubleshoot this issue.


Figure 7-18 *Troubleshooting Flowchart for Problems Related to EIGRP Not Advertising Routes to Its Neighbors*



EIGRP Is Not Advertising Routes to Its Neighbors—Cause: Distribute List

Figure 7-19 shows a network in which EIGRP is not advertising routes to its neighbor because of a distribute list problem. Example 7-14 shows the configurations for Routers A and B in this network.

Figure 7-19 *EIGRP Network Not Advertising Routes to Its Neighbors Because of a Misconfigured Distribute List*



Example 7-14 Configurations for Routers A and B in Figure 7-19

```

Router A# interface ethernet 0
    ip address 172.16.3.1 255.255.255.0
interface serial 0
    ip address 10.1.1.1 255.255.255.0
router eigrp 1
    network 172.16.0.0
    network 10.0.0.0

Router B# interface ethernet 0
    ip address 192.168.3.17 255.255.255.240
interface serial 0
    ip address 10.1.1.2 255.255.255.0
router eigrp 1
    network 192.168.3.0
    network 10.0.0.0
    distribute-list 1 out
    access-list 1 permit 192.168.3.160 0.0.0.15

```

The problem is that Router A is not receiving the routes from Router B about network 192.168.3.16. Example 7-15 shows the debug output on Router B.

Example 7-15 debug ip eigrp Command Output on Router B

```

Router_B# debug ip eigrp

IP-EIGRP: 192.168.3.16/28 - denied by distribute list

```

As the output in Example 7-15 reveals, Router B won't advertise the 192.168.3.16 because of the distribute list configuration. Looking again at the configuration in Example 7-14, you can see that the distribute list is tied to access-list 1, and access-list 1 has the network number misconfigured. access-list 1 should permit 192.168.3.16 instead of 192.168.3.160. Because 192.168.3.16 is not included in the **permit** statement, there is an implicit **deny** in the access list that prevents network 192.168.3.16 being advertised.

The solution to this problem is to change access-list 1 to permit 192.168.3.16 instead of 192.168.3.160. Changing the access list to permit 192.168.3.16 fixes the problem.

EIGRP Is Not Advertising Routes to Its Neighbors—Cause: Discontiguous Networks

Using the network diagram in Figure 7-19, another issue with EIGRP not advertising the network could be manual summarization configured on the interface or autosummarization across a major network boundary, as shown in Example 7-16.

Example 7-16 Configurations for Routers A and B in Figure 7-19

```

Router A# interface ethernet 0
    ip address 192.168.3.33 255.255.255.240
interface serial 0
    ip address 10.1.1.1 255.255.255.0
router eigrp 1
    network 192.168.3.0
    network 10.0.0.0

```

Example 7-16 *Configurations for Routers A and B in Figure 7-19 (Continued)*

```

Router B# interface ethernet 0
      ip address 192.168.3.21 255.255.255.240
interface serial 0
      ip address 10.1.1.2 255.255.255.0
router eigrp 1
      network 192.168.3.0
      network 10.0.0.0

```

The problem is that Router A is not receiving routes for the 192.168.3.16 network from Router B. Example 7-17 shows the debug output on Router B.

Example 7-17 *debug ip eigrp Command Output on Router B*

```

Router B# debug ip eigrp

IP-EIGRP: 192.168.3.16/28 -don't advertise out Serial0
IP-EIGRP: 192.168.3.0/24 - do advertise out Serial0

```

From the debug, Router B shows that it is not advertising the 192.168.3.16/28 network; however, it is advertising only the major network of 192.168.3.0/24 to Router A. Looking at the configuration of Routers A and B in Example 7-16 shows that the two routers have a discontinuous network. Router A has the network of 192.168.3.32/28 in its Ethernet, while Router B has another network of 192.168.3.16/28 in its Ethernet, separated by a network of 10.1.1.0/24. Therefore, when Router B advertises the network of 192.168.3.16/28 across a major network boundary of 10.1.1.0, it advertises only the major network of 192.168.3.0/24 to Router A instead of advertising the network of 192.168.3.16/28. When Router A receives the major network of 192.168.3.0/24, it does not install the network in the topology table because it already has the 192.168.3.0 network on its Ethernet interface.

Two solutions to the discontinuous network problem exist. One is to configure the command **no auto-summary** under **router eigrp**. This command tells EIGRP not to autosummarize to major network boundaries. As a result, Router B's configuration will look like Example 7-18.

Example 7-18 *Disabling Autosummarization on Router B to Prevent Discontinuous Networks*

```

Router B# router EIGRP 1
      network 192.168.3.0
      network 10.0.0.0
      no auto-summary

```

The second solution is to change the IP address of the serial interfaces on each side of the link to the 192.168.3.0 subnet. As an example, the serial IP address can take 192.168.3.65/28 and 192.168.3.66/28. This way, Router B won't autosummarize the route because it is not across a major network boundary.


EIGRP Is Not Advertising Routes to Neighbors—Cause: Split-Horizon Issues

EIGRP has its own **split-horizon** command. This command, configured under the interface, is shown here:

```
[no] ip split-horizon eigrp autonomous-system
```

Turning off IP split horizon does not turn off EIGRP split horizon. Figure 7-20 shows an EIGRP network vulnerable to split-horizon issues.

Figure 7-20 *EIGRP Network Susceptible to EIGRP Split-Horizon Problems*



Example 7-19 shows the configurations for Routers A, B, and C in the hub-and-spoke network in Figure 7-20.

Example 7-19 *Configurations for Routers A, B, and C in Figure 7-20*

```

Router A# interface ethernet 0
      ip address 172.16.1.1 255.255.255.0
interface serial 0
      ip address 172.16.2.1 255.255.255.0
router eigrp 1
  network 172.16.0.0

```

```

Router B# interface ethernet 0
      ip address 172.16.3.1 255.255.255.0
interface serial 0
      ip address 172.16.2.2 255.255.255.0
router eigrp 1
  network 172.16.0.0

```

```

Router C# interface ethernet 0
      ip address 172.16.4.1 255.255.255.0
interface serial 0
      ip address 172.16.2.3 255.255.255.0
router eigrp 1
  network 172.16.0.0

```

A common network environment, shown in Figure 7-20 is the Frame Relay hub-and-spoke design, in which the hub router (Router A) in Figure 7-20 doesn't have a subinterface configured

for each remote spoke site. As a result, the hub router uses a main interface to connect to the two spoke sites. The problem is that Router B doesn't receive the routes for Router C's Ethernet network of 172.16.4.0/24, and Router C doesn't receive the routes for Router B's Ethernet network of 172.16.3.0/24. The problem seems to be at the hub site. The hub site sees all the routes, but the hub site is not passing the routes from Router B to Router C, and vice versa. Example 7-20 shows the **debug** output on the hub router (Router A).

Example 7-20 *debug ip eigrp Command Output on Router A*

```
Router A# debug ip eigrp
IP-EIGRP: 172.16.1.0/24 - do advertise out Serial0
IP-EIGRP: Processing incoming UPDATE packet
IP-EIGRP: Int 172.16.3.0/24
IP-EIGRP: Int 172.16.4.0/24
```

From the debug, you can see that the hub router advertises only the 172.16.1.0/24 route on Serial0. The hub router receives routes for the 172.16.3.0/24 and 172.16.4.0/24 interfaces from Router B and Router C. The problem is that the hub router is not sending all the routes on Serial0. Referring to the configurations of Routers A, B, and C in Example 7-19, you can see that their serial interfaces are all in the same subnet, but they are not physically connected. Therefore, the hub router receives the routes from Serial0 from Router B and Router C but won't readvertise those routes on Serial0. This follows the split-horizon rule (route information must not exit the router interface through which that information was received).

To solve the split-horizon problem for EIGRP, the easiest fix is to turn off split horizon for EIGRP. Example 7-21 shows the correct configuration change to disable split horizon.

Example 7-21 *Disabling Split Horizon on the Hub Router*

```
Router A# interface ethernet 0
ip address 172.16.1.1 255.255.255.0
interface serial 0
ip address 172.16.2.1 255.255.255.0
no IP split-horizon EIGRP 1
router EIGRP 1
network 172.16.0.0
```

Example 7-22 shows the debug output on Router A after the configuration change.

Example 7-22 *Verifying That Disabling Split Horizon Corrected the Problem*

```
Router A# debug ip eigrp
IP-EIGRP: 172.16.1.0/24 - do advertise out Serial0
IP-EIGRP: 172.16.3.0/24 - do advertise out Serial0
IP-EIGRP: 172.16.4.0/24 - do advertise out Serial0
IP-EIGRP: Processing incoming UPDATE packet
IP-EIGRP: Int 172.16.3.0/24
IP-EIGRP: Int 172.16.4.0/24
```

Now the spoke Routers B and C can see the routes. Another fix for the split-horizon problem is to configure subinterfaces on the hub router and assign different IP address subnets for each subinterface. Keep in mind that the support of a serial subinterface is

valid for only the WAN PVC type of connection, such as ATM or Frame Relay. Example 7-23 shows the configuration for such a setup to avoid the EIGRP split-horizon problem.

Example 7-23 *Configuring Subinterfaces with Different IP Address Subnets to Combat EIGRP Split-Horizon Problems*

```
Router A# interface ethernet 0
ip address 172.16.1.1 255.255.255.0
interface serial 0.1 point-to-point
description connection to router B
ip address 172.16.2.1 255.255.255.0
interface serial 0.2 point-to-point
description connection to router C
ip address 172.16.5.1 255.255.255.0
router eigrp 1
network 172.16.0.0
```

```
Router B# interface ethernet 0
ip address 172.16.3.1 255.255.255.0
interface serial 0
ip address 172.16.2.2 255.255.255.0
router eigrp 1
network 172.16.0.0
```

```
Router C# interface ethernet 0
ip address 172.16.4.1 255.255.255.0
interface serial 0
ip address 172.16.5.2 255.255.255.0
router eigrp 1
network 172.16.0.0
```

When subinterfaces are configured in Router A, this logically separates the connection to Router B and Router C. Each connection to Router B and Router C has its own network. For example, the connection from Router A to Router B is now through connection Serial 0.1 over the 172.16.2.0/24 network, and the connection from Router A to Router C is now through connection Serial 0.2 over the 172.16.5.0/24 network. Because Router A has two logical connection to Routers B and C over two different logical interfaces, the split horizon rule doesn't apply and Router A will advertise all the routes to routers B and C, as shown in Example 7-24.

Example 7-24 *Verifying That Configuring the Subinterface with Different Subnets Solves the Split-Horizon Problem*

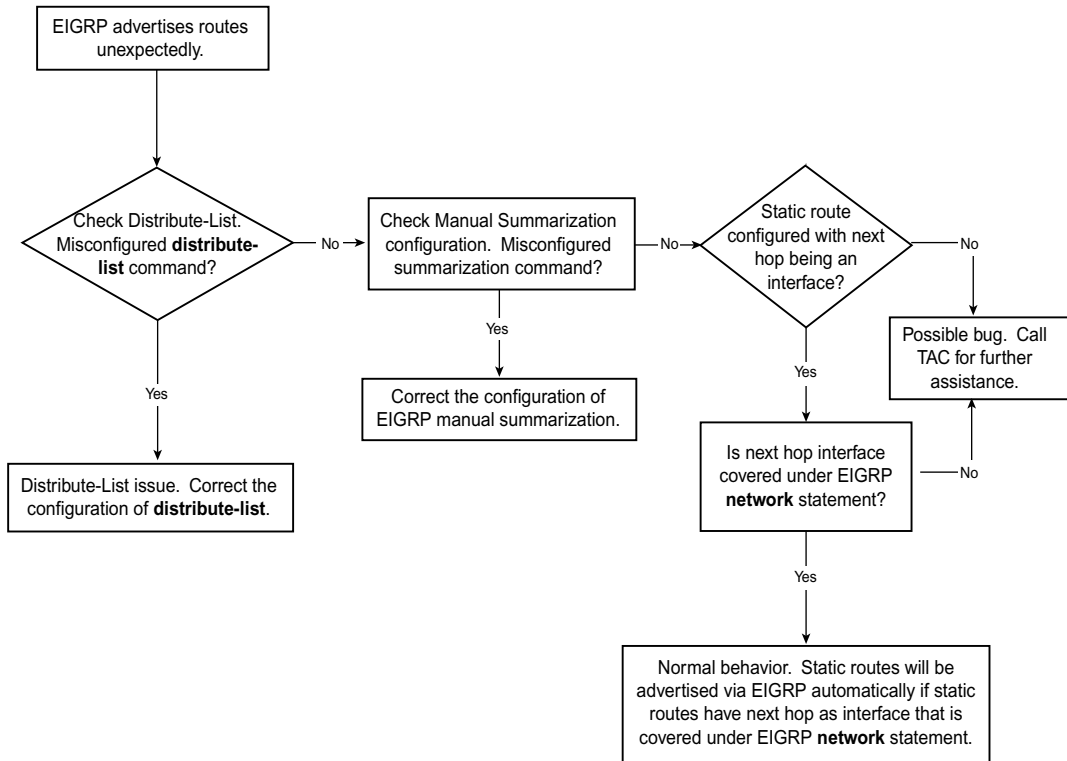
```
Router A# debug ip eigrp
IP-EIGRP: 172.16.1.0/24 - do advertise out Serial0.1
IP-EIGRP: 172.16.4.0/24 - do advertise out Serial0.1
IP-EIGRP: 172.16.5.0/24 - do advertise out Serial0.1
IP-EIGRP: 172.16.1.0/24 - do advertise out Serial0.2
IP-EIGRP: 172.16.2.0/24 - do advertise out Serial0.2
IP-EIGRP: 172.16.3.0/24 - do advertise out Serial0.2
```

With Router A advertising all the routes to the remote Routers, Routers B and C now can reach each other's LAN interface.

EIGRP Is Advertising Routes to Neighbors When the Network Administrators Think That It Shouldn't

Sometimes, EIGRP advertises unexpected routes to its neighbors. See Figure 7–21 for a flowchart of troubleshooting EIGRP unexpected advertisement of routes.

Figure 7-21 Flowchart for Troubleshooting EIGRP Unexpected Advertisement of Routes



Refer to Figure 7-19 for the network diagram on this example. Example 7-25 shows the configurations for Routers A and B.

Example 7-25 Configuration of Router A and Router B for the Example Shown in Figure 7-19

```

Router A# interface ethernet 0
ip address 172.16.3.1 255.255.255.0
interface serial 0
ip address 10.1.1.1 255.255.255.0
router eigrp 1
network 172.16.0.0
network 10.0.0.0
    
```

```

Router B# interface ethernet 0
ip address 192.168.130.1 255.255.255.0
    
```

continues

Example 7-25 Configuration of Router A and Router B for the Example Shown in Figure 7-19 (Continued)

```

interface serial 0
  ip address 10.1.1.2 255.255.255.0
router eigrp 1
  network 192.168.130.0
  network 10.0.0.0
ip route 192.168.1.0 255.255.255.0 ethernet 0
ip route 192.168.2.0 255.255.255.0 ethernet 0
ip route 192.168.3.0 255.255.255.0 ethernet 0
ip route 192.168.4.0 255.255.255.0 ethernet 0
.
.
.
ip route 192.168.127.0 255.255.255.0 ethernet 0

```

The problem is that, without inserting the **redistribute static** command under the **router eigrp** command in Router B, Router B automatically redistributes all the 127 static routes configured to Router A. This can cause unnecessary routes being advertised inadvertently throughout the entire network. The cause of the problem is that the static routes are configured with the outbound interface. In this case, the router thinks that all the static routes are directly connected to the Ethernet 0 interface. These Ethernet interfaces also are covered under the router EIGRP process by the **network 192.168.130.0** command. Because Ethernet 0 is considered to run EIGRP, all the networks connected to it by a static route also are considered to belong to the EIGRP process. The router then advertises all these static routes even though **redistribute static** is not configured.

The solution to this problem is either to configure a distribute list that prevents the router from advertising all those static routes or to change the static routes to reference the next-hop IP addresses instead of an interface. This way, the router will not advertise all these static routes and flood the entire network with unnecessary routes.

Example 7-26 shows the distribute list configured on Router B to stop sending the unwanted redistributed static routes.

Example 7-26 Configuration on Router B to Stop Sending Unwanted Static Routes by Configuring Distribute List

```

Router B# interface ethernet 0
  ip address 192.168.130.1 255.255.255.0
interface serial 0
  ip address 10.1.1.2 255.255.255.0
router eigrp 1
  network 192.168.130.0
  network 10.0.0.0
  distribute-list 1 out
ip route 192.168.1.0 255.255.255.0 ethernet 0
ip route 192.168.2.0 255.255.255.0 ethernet 0
ip route 192.168.3.0 255.255.255.0 ethernet 0
ip route 192.168.4.0 255.255.255.0 ethernet 0
.
.
.
ip route 192.168.127.0 255.255.255.0 ethernet 0
access-list 1 deny 192.168.0.0 0.0.127.255
access-list 1 permit any

```

The distribute list is tied to access-list 1, and access-list 1 denies sending out any routes that ranges from 192.168.0.0/24 through 192.168.127.0/24 and permits sending any other routes. Such a distribute list stops sending out the unwanted redistributed static routes in the example. The debug output on Router B, shown in Example 7-27, shows that the router does not send the static routes to other EIGRP neighbors because the distribute list is configured.

Example 7-27 *Verification on Router B Not Sending Out Static Routes Because a Distribute List Is Configured*

```
Router B# debug ip eigrp
IP-EIGRP: 192.168.1.0/24 - denied by distribute list
IP-EIGRP: 192.168.2.0/24 - denied by distribute list
IP-EIGRP: 192.168.3.0/24 - denied by distribute list
IP-EIGRP: 192.168.4.0/24 - denied by distribute list
IP-EIGRP: 192.168.5.0/24 - denied by distribute list
IP-EIGRP: 192.168.6.0/24 - denied by distribute list
.
.
.
IP-EIGRP: 192.168.127.0/24 - denied by distribute list
```

The other solution to this problem is to redefine the static routes so that the next hop of the static route is an IP address instead of an interface. Example 7-28 shows the change of static route configuration in Router B to fix the problem.

Example 7-28 *Configuration on Router B to Stop Sending Unwanted Static Routes by Reconfiguring Static Routes with the Next Hop—an IP Address Instead of an Interface*

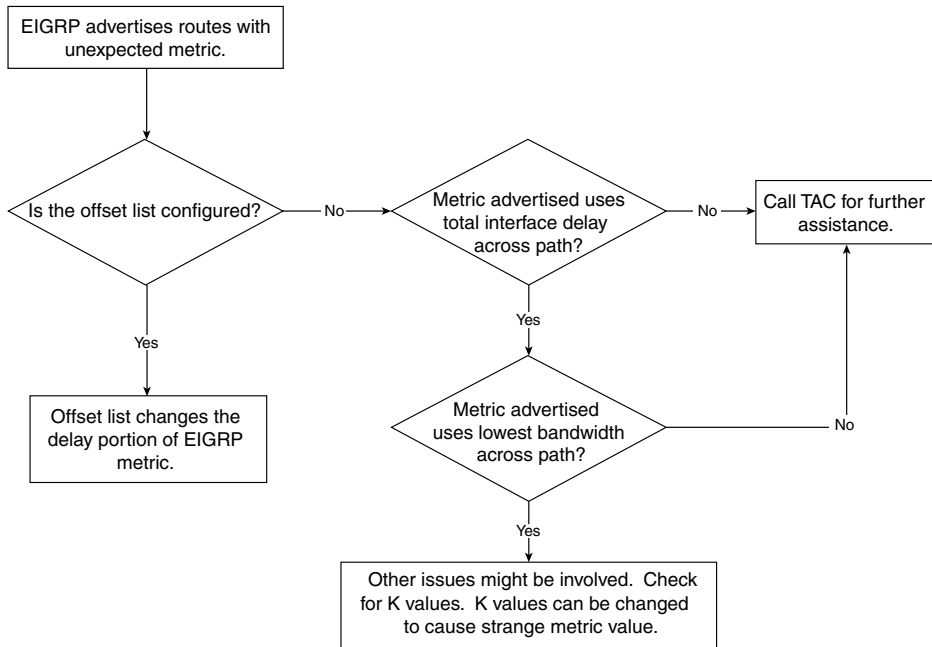
```
Router B# interface ethernet 0
ip address 192.168.130.1 255.255.255.0
interface serial 0
ip address 10.1.1.2 255.255.255.0
router eigrp 1
network 192.168.130.0
network 10.0.0.0
distribute-list 1 out
ip route 192.168.1.0 255.255.255.0 192.168.130.2
ip route 192.168.2.0 255.255.255.0 192.168.130.2
ip route 192.168.3.0 255.255.255.0 192.168.130.2
ip route 192.168.4.0 255.255.255.0 192.168.130.2
.
.
.
ip route 192.168.127.0 255.255.255.0 192.168.130.2
```

EIGRP Is Advertising Routes with Unexpected Metric

Not only might EIGRP advertise unexpected routes to its neighbors, but it also might advertise an unexpected metric to its neighbors. The EIGRP metric is the basis of route selection done by EIGRP, which selects the route with the lowest EIGRP metric to the destination network. An unexpected EIGRP metric being sent or received on the router


might alter route selection to the destination network. The end result might be suboptimal routing. Figure 7-22 shows the flowchart for troubleshooting such an issue.

Figure 7-22 Flowchart for Troubleshooting EIGRP Advertisement of Routes with Unexpected Metric Value



The case study that follows is a case of an offset list that is created inadvertently, causing the router to route packets in a suboptimal fashion. The **offset-list** command adds an offset value to the routing metrics. It's a way to manipulate the routing metric for certain routes, thereby, altering the route selection for a particular routing protocol. Figure 7-23 illustrates the network setup for the unexpected metric value problem.

Figure 7-23 EIGRP Network Susceptible to EIGRP Advertisement Problems Because of Unexpected Metric Values



Example 7-29 shows the configurations for the routers in the EIGRP network shown in Figure 7-23.

Example 7-29 Configurations for Routers A, B, and C in Figure 7-23

<pre> Router A# interface ethernet 0 ip address 172.16.1.1 255.255.255.0 interface serial 0 ip address 172.16.2.1 255.255.255.0 interface serial 1 ip address 172.16.3.1 255.255.255.0 router eigrp 1 network 172.16.0.0 </pre>
<pre> Router B# interface ethernet 0 ip address 172.16.6.1 255.255.255.0 interface serial 0 ip address 172.16.2.2 255.255.255.0 interface serial 1 ip address 172.16.4.1 255.255.255.0 router eigrp 1 network 172.16.0.0 </pre>
<pre> Router C# interface ethernet 0 ip address 172.16.5.1 255.255.255.0 interface serial 0 ip address 172.16.4.2 255.255.255.0 interface serial 1 ip address 172.16.3.2 255.255.255.0 router eigrp 1 network 172.16.0.0 offset-list 1 out 6000000 serial 1 access-list 1 permit 172.16.0.0 0.0.255.255 </pre>

The problem is that Router A is not taking the direct paths to Router C to reach Router C's Ethernet network of 172.16.5.0/24. Instead, Router A takes the path to Router B and then to Router C. This takes an extra hop. Example 7-30 shows the routing table and the EIGRP topology table for 172.16.5.0 255.255.255.0 for Router A.

Example 7-30 `show ip route` and `show ip eigrp topology` Command Output Reveals the Routes That Router A Is Taking to Reach Router C's 172.16.5.0/24 Ethernet Network

<pre> Router_A#show ip route 172.16.5.0 Routing entry for 172.16.5.0/24 Known via "eigrp 1", distance 90, metric 2707456, type internal Redistributing via eigrp 1 Last update from 172.16.2.2 on Serial0, 01:08:13 ago Routing Descriptor Blocks: *172.16.2.2, from 172.16.2.2, 01:08:13 ago, via Serial0 Route metric is 2707456, traffic share count is 1 Total delay is 41000 microseconds, minimum bandwidth is 1544 Kbit Reliability 255/255, minimum MTU 1500 bytes Loading 1/255, Hops 2 </pre>
<pre> Router A# show ip eigrp topology 172.16.5.0 255.255.255.0 IP-EIGRP topology </pre>

continues

Example 7-30 *show ip route and show ip eigrp topology Command Output Reveals the Routes That Router A Is Taking to Reach Router C's 172.16.5.0/24 Ethernet Network (Continued)*

```

entry for 172.16.5.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2707456
Routing Descriptor Blocks:
 172.16.2.2 (Serial0), from 172.16.2.2, Send flag is 0x0
    Composite metric is (2707456/2195456), Route is Internal
    Vector metric:
      Minimum bandwidth is 1544 Kbit
      Total delay is 41000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 2
 172.16.3.2 (Serial1), from 172.16.3.2, Send flag is 0x0
    Composite metric is (2795456/281600), Route is Internal
    Vector metric:
      Minimum bandwidth is 1544 Kbit
      Total delay is 44437 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 1

```

Example 7-30 shows that Router A chooses Router B as the next hop to Router C because Router B has a better metric than Router C. Looking in detail at the topology table shows that the path to Router C has more delay than the path to Router B, but all the links are T1 links. The interface configuration in Router C didn't show any manually configured delay value. Looking at the configuration in Router C more in detail reveals the offset-list configuration under **router eigrp** in Router C.

The offset list in Router C adds a metric of 600,000 to outgoing routes in Serial1. This is the cause of the problem. The offset values added increase the delay value when Router C sends the routes to Router A, causing Router A to prefer routes from Router B.

The solution is to remove the offset list configured on Router C. To remove the offset list, configure Router C as in Example 7-31.

Example 7-31 *Removing the Offset List from Router C's Configuration*

```

Router C# config term
Router_C(config)#router eigrp 1
Router_C(config-router)#no offset-list 1 out 600000 serial 1

```

Example 7-32 shows the routing table and the topology table in Router A after removing the offset list configured on Router C.

Example 7-32 *show ip route and show ip eigrp topology Command Output Verifies That Router A Is Now Taking the Optimal Routes to Reach Router C's 172.16.5.0/24 Ethernet Network*

Router_A#show ip route 172.16.5.0
Routing entry for 172.16.5.0/24
Known via "eigrp 1", distance 90, metric 2195456, type internal
Redistributing via eigrp 1
Last update from 172.16.3.2 on Serial1, 00:08:23 ago
Routing Descriptor Blocks:
*172.16.3.2, from 172.16.3.2, 00:08:23 ago, via Serial1
Route metric is 2195456, traffic share count is 1
Total delay is 21000 microseconds, minimum bandwidth is 1544 Kbit
Reliability 255/255, minimum MTU 1500 bytes
Loading 1/255, Hops 1
Router A# show ip eigrp topology 172.16.5.0 255.255.255.0
IP-EIGRP topology entry for 172.16.5.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2195456
Routing Descriptor Blocks:
172.16.3.2 (Serial1), from 172.16.3.2, Send flag is 0x0
Composite metric is (2195456/281600), Route is Internal
Vector metric:
Minimum bandwidth is 1544 Kbit
Total delay is 21000 microseconds
Reliability is 255/255
Load is 1/255
Minimum MTU is 1500
Hop count is 1
172.16.2.2 (Serial1), from 172.16.2.2, Send flag is 0x0
Composite metric is (2707456/2195456), Route is Internal
Vector metric:
Minimum bandwidth is 1544 Kbit
Total delay is 41000 microseconds
Reliability is 255/255
Load is 1/255
Minimum MTU is 1500
Hop count is 2

The output in Example 7-32 now shows 172.16.3.2 as the next hop to Router C, which is the optimal path to the 172.16.5.0/24 network. Also, compare the topology table shown in Example 7-30 and Example 7-32. The EIGRP metric coming from the neighbor 172.16.3.2 has been reduced from the metric of 2,795,456 to 2,195,456. This reduction of metric of 600,000 is the result of removing the offset list. As this case study demonstrates, it is important that you scrutinize the configuration when abnormal behavior occurs. When opening a case with Cisco's TAC, be sure to provide router configuration whenever possible.

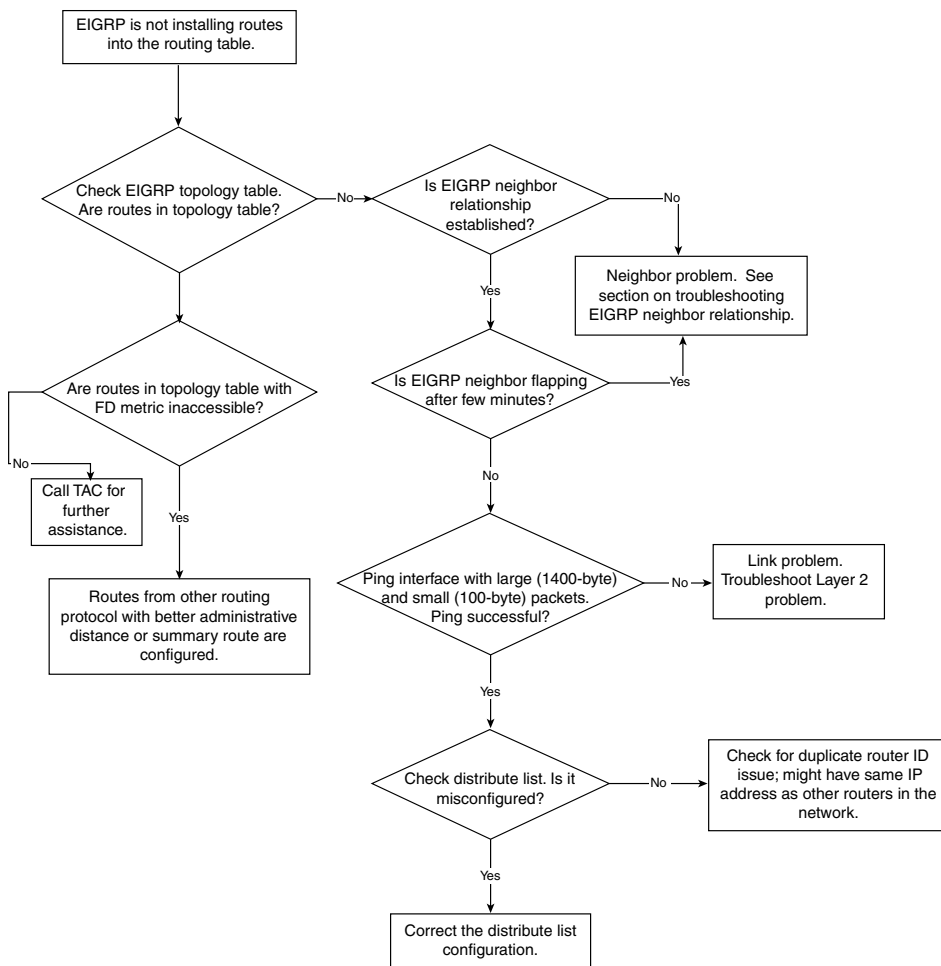
Troubleshooting EIGRP Route Installation

The previous section discusses the problems that EIGRP routers have when advertising routes to its neighbors. This section discusses troubleshooting problems when EIGRP doesn't install the routes in the routing table. The most common causes of this problem are as follows:

- Auto or manual summarization configured
- Higher administrative distance
- Duplicate router IDs

The following sections detail the causes of this problem and how to resolve them. For overall troubleshooting methods, Figure 7-24 shows the flowchart for troubleshooting EIGRP route-installation problems.


Figure 7-24 Flowchart for Troubleshooting EIGRP Route-Installation Problems



EIGRP Is Not Installing Routes—Cause: Auto or Manual Summarization

When EIGRP fails to install routes in the routing table, the first thing to check is the topology table. Figure 7-25 shows the network setup for this case study.

Figure 7-25 EIGRP Network Susceptible to Route-Installation Problem



Example 7-33 shows the configuration for Router B.

Example 7-33 Configuration for Router B in Figure 7-25

```

Router B# interface ethernet 0
ip address 172.16.2.1 255.255.255.0
interface serial 0
192.168.1.1 255.255.255.0
interface serial 1
192.168.2.1 255.255.255.0
router eigrp 1
network 172.16.0.0
network 192.168.1.0
network 192.168.2.0
  
```

Inside network clouds X and Y are networks in the 172.16.x.x space. The problem is that Router C summarizes all the 172.16.x.x networks into one summary route of 172.16.0.0/16 and sends it to Router B. Router B is not installing the routes in the routing table, as shown in Example 7-34.

Example 7-34 Router B's Routing Table

```

Router B# show ip route 172.16.0.0

Routing entry for 172.16.0.0/16
  Routing Descriptor Blocks:
    * directly connected, via Null 0
  
```


Router B's routing table shows that the route is directly connected to Null 0 instead of learned from Router C. The topology table in Router B shows that the router is getting the routes from Router C but is installing the route as connected because the Null 0 route has a distance of 5, which is an EIGRP summary route. The configuration of Router B shows that EIGRP summarizes the 172.16.0.0/16 route because of autosummarization. Every time autosummarization or manual summarization takes place, EIGRP installs the summary route with the next hop to Null 0. This is a loop-prevention mechanism for EIGRP's summary routes. In this case study, this is exactly what happens—EIGRP does not install a route from its neighbor that falls within its summary range.

The solution to this problem, based on this cause, is more of a design issue. Two places in the network must not send the same summary routes to one another. In this example, you configure the **no auto-summary** command on Router B to allow Router B to accept the summary routes coming from Router C. Example 7-35 shows the configuration in Router B to fix the problem.

Example 7-35 Configuration Change on Router B to Fix the Problem Shown in Figure 7-25

```
Router B# interface ethernet 0
      ip address 172.16.2.1 255.255.255.0
interface serial 0
      192.168.1.1 255.255.255.0
interface serial 1
      192.168.2.1 255.255.255.0
router eigrp 1
      network 172.16.0.0
      network 192.168.1.0
      network 192.168.2.0
      no auto-summary
```

With the configuration change in Router B, the routing table shown in Example 7-36 for Router B now shows the summary route of 172.16.0.0/16 coming from Router C.

Example 7-36 Routing Table of Router B Now Showing Summary Route Coming from Router C

```
Router_B#show ip route 172.16.0.0 255.255.0.0
Routing entry for 172.16.0.0/16
  Known via "eigrp 1", distance 90, metric 2195456, type internal
  Redistributing via eigrp 1
  Last update from 192.168.1.2 on Serial0, 00:16:24 ago
  Routing Descriptor Blocks:
    *192.168.1.2, from 192.168.1.2, 00:16:24 ago, via Serial0
      Route metric is 2195456, traffic share count is 1
      Total delay is 21000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

EIGRP Is Not Installing Routes—Cause: Higher Administrative Distance

Refer to the network topology in Figure 7-25. Another variation of a similar problem can happen if network cloud Y sends external EIGRP routes of 150.150.0.0/16 to Router B, and Router B is running RIP and EIGRP but is getting the 150.150.0.0/16 routes from the RIP domain from Router A. Because RIP has a lower administrative distance (120) than external EIGRP routes (170), Router B installs RIP routes for 150.150.0.0/16 only from Router A. Example 7-37 shows the EIGRP topology table for Router B.

Example 7-37 Router B's EIGRP Topology Table for 150.150.0.0/16

```
Router B# show ip eigrp topology 150.150.0.0 255.255.0.0

IP-EIGRP topology entry for 150.150.0.0/16
State is Passive, Query origin flag is 1, 0 Successor(s), FD is 4294967295
Routing Descriptor Blocks:
  192.168.1.2 (Serial0), from 192.168.1.2, Send flag is 0x0
  Composite metric is (2707456/2195456), Route is External
    Vector metric:
      Minimum bandwidth is 1544 Kbit
      Total delay is 41000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 3
    External data:
      Originating router is 155.155.155.1
      AS number of routes is 0
      External protocol is OSPF, external metric is 64
      Administrator tag is 0
```

The EIGRP topology table shows that the feasible distance (FD) is inaccessible (4294967295); the route is an external route that has been redistributed from OSPF. This means that Router B is receiving the 150.150.0.0/16 routes from Router C but is setting the FD as inaccessible because Router B is *not* using the EIGRP route in the routing table. As a matter of fact, the routing table entry in Router B is a RIP route for 150.150.0.0/16, as shown in Example 7-38. In other words, when the FD is inaccessible in the EIGRP topology table, the router is not using that EIGRP route in its routing table. Usually, the route is overridden by another routing protocol that has lower administrative distance.

Example 7-38 Routing Table of Router B Showing 150.150.0.0/16 Route as a RIP Route

```
Router_B#show ip route 150.150.0.0
Routing entry for 150.150.0.0/16
  Known via "rip", distance 120, metric 5
  Redistributing via rip
  Last update from 192.168.2.2 on Serial1, 00:00:24 ago
  Routing Descriptor Blocks:
    *192.168.2.2, from 192.168.2.2, 00:00:24 ago, via Serial1
      Route metric is 5, traffic share count is 1
```

To fix this problem, you must change the administrative distance of the routing protocols so that external EIGRP routes are preferred. To do so, use the **distance** command to manipulate the administrative distance of a routing protocol. The configuration of Router B to fix this problem is shown in Example 7-39.

Example 7-39 *Configuration Change on Router B to Fix the Route-Installation Problem Because of Higher Administrative Distance*

```
Router B# interface ethernet 0
      ip address 172.16.2.1 255.255.255.0
      interface serial 0
        192.168.1.1 255.255.255.0
      interface serial 1
        192.168.2.1 255.255.255.0
      router eigrp 1
        network 172.16.0.0
        network 192.168.1.0
        network 192.168.2.0
      router rip
        network 172.16.0.0
        network 192.168.2.0
        distance 180 192.168.2.2 255.255.255.255
```

The **distance** command shown in Example 7-39 sets the RIP administrative distance to 180 for any updates coming from 192.168.2.2. This allows the external EIGRP routes (administrative distance of 170) coming from Router C to be preferred over RIP routes. Example 7-40 shows the result.


Example 7-40 *Routing Table of Router B Now Showing Summary Route Coming from Router C*

```
Router_B#show ip route 150.150.0.0
Routing entry for 150.150.0.0/16
  Known via "eigrp 1", distance 90, metric 2195456, type internal
  Redistributing via eigrp 1
  Last update from 192.168.1.2 on Serial0, 00:26:14 ago
  Routing Descriptor Blocks:
    *192.168.1.2, from 192.168.1.2, 00:26:14 ago, via Serial0
      Route metric is 2195456, traffic share count is 1
      Total delay is 21000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

EIGRP Is Not Installing Routes—Cause: Duplicate Router IDs

Many times, EIGRP will not install routes because of a duplicate router ID problem. EIGRP does not use router ID as extensively as OSPF. EIGRP uses the notion of router ID only on external routes to prevent loops. EIGRP chooses the router ID based on the highest IP address of the loopback interfaces on the router. If the router doesn't have any loopback interfaces, the highest active IP address of all the interfaces is chosen as the router ID for EIGRP. Figure 7-26 shows the network setup for such a case study on EIGRP router IDs.

Figure 7-26 EIGRP Network Susceptible to EIGRP Not Installing Routes Because of Duplicate Router IDs



Example 7-41 shows the pertinent configurations for the cause of this problem.

Example 7-41 Configurations for Routers A, B, C, and X in Figure 7-26

<pre>Router A# interface ethernet 0 ip address 192.168.1.1 255.255.255.0 interface serial 0 ip address 10.1.1.1 255.255.255.0</pre>
<pre>Router B# interface serial 0 IP address 10.1.1.2 255.255.255.0 interface serial 1 IP address 10.1.2.1 255.255.255.0</pre>
<pre>Router C# interface serial 0 ip address 10.1.2.2 255.255.255.0</pre>
<pre>Router X# interface loopback 0 ip address 192.168.1.1 255.255.255.255</pre>

Router X is redistributing a route of 150.150.0.0/16 from OSPF into EIGRP and is sending the route several hops to Router C. Router C receives the route and sends the route as EIGRP external routes to Router B. Router B installs the route in the routing table and sends it to Router A. The debug output in Example 7-42 verifies how Router B sends the route to Router A.

Example 7-42 debug ip eigrp Command Output on Router B

<pre>Router B# debug ip eigrp IP-EIGRP: 150.150.0.0/16 - do advertise out serial 0</pre>
--

The problem is that Router A is not installing the 150.150.0.0/16 route in the routing table. As a matter of fact, Router A is not showing the 150.150.0.0/16 route in its topology table. Going back to Router B, the route is in the routing table, and the topology table appears as shown in Example 7-43.

Example 7-43 *EIGRP Topology Table for 150.150.0.0/16 on Router B*

Router B# show ip eigrp topology 150.150.0.0 255.255.0.0
IP-EIGRP topology entry for 150.150.0.0/16
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 3757056
Routing Descriptor Blocks:
10.1.2.2 (Serial1), from 10.1.2.2, Send flag is 0x0
Composite metric is (3757056/3245056), Route is External
Vector metric:
Minimum bandwidth is 1544 Kbit
Total delay is 82000 microseconds
Reliability is 255/255
Load is 1/255
Minimum MTU is 1500
Hop count is 7
External data:
Originating router is 192.168.1.1
AS number of routes is 0
External protocol is OSPF, external metric is 64
Administrator tag is 0

Router B shows that it is getting the routes from Router C. By looking at the external data section, notice that the originating router is 192.168.1.1, which is seven hops away. The original protocol that originated the route 150.150.0.0/16 is OSPF with the metric of 64. Notice that the originating router is 192.168.1.1. Looking back at the configuration of Router A in Example 7-41, notice that Router A also has an IP address of 192.168.1.1 configured on Ethernet 0, and it is the highest IP address on the router. All this evidence points to a duplicate router ID problem in EIGRP that causes Router A not to install routes. Because Router X and Router A have the same router ID (192.168.1.1), when Router A receives the route from Router B, it looks at the external data section of the route to see who is the originating router. In this case, Router A sees the originating router as 192.168.1.1, which is its own router ID. Router A does not put the route in its topology table because it thinks that it is the originator of the route and that by receiving the route back from other neighbors, it must be a loop. So, to prevent a routing loop, Router A does not put the route of 150.150.0.0/16 in the topology table. Consequently, the route does not appear in the routing table.

Router A will not install any external routes that originate from Router X because external routes carry the router ID in their EIGRP update packet. Router A will install internal EIGRP routes from Router X without any problem. The duplicate router ID problem happens only for external routes.

The solution to the duplicate router ID problem is to change the IP address of the loopback interface of Router X or to change the IP address of Ethernet 0 in Router A. The rule of thumb: Never configure the same IP address on two places in the network. Change the loopback IP address of Router X to 192.168.9.1/32 to fix this problem (see Example 7-44).

The result of the IP address change in Router X is the installment of the 150.150.0.0/16 route in Router A, as shown in Example 7-45.

Example 7-44 *Loopback IP Address Change in Router X to Avoid Duplicate Router ID Problem*

```
Router X#interface Loopback 0
IP address 192.168.9.1 255.255.255.255
```

Example 7-45 *Routing Table and EIGRP Topology Table for 150.150.0.0/16 on Router A to Verify the Fix*

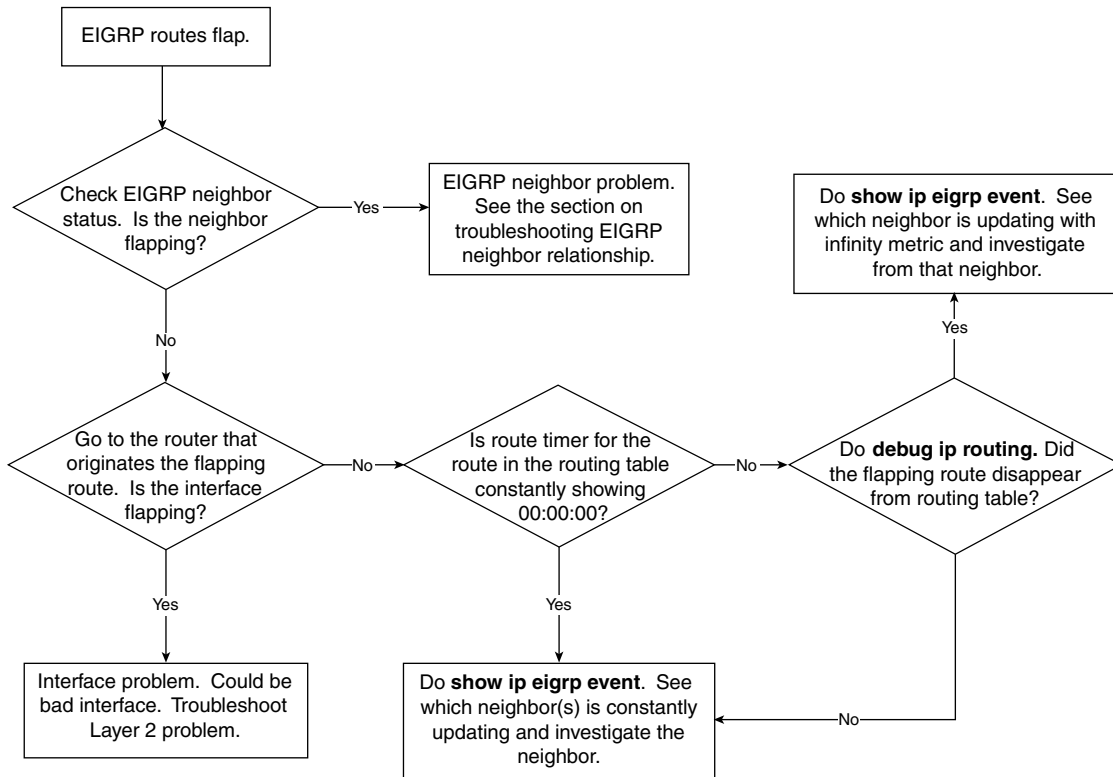
```
Router_A#show ip route 150.150.0.0
Routing entry for 150.150.0.0/16
  Known via "eigrp 1", distance 170, metric 4269056, type external
  Redistributing via eigrp 1
  Last update from 10.1.1.2 on Serial0, 00:06:14 ago
  Routing Descriptor Blocks:
    *10.1.1.2, from 10.1.1.2, 00:06:14 ago, via Serial0
      Route metric is 4269056, traffic share count is 1
      Total delay is 102000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 8

Router A# show ip eigrp topology 150.150.0.0 255.255.0.0
IP-EIGRP topology entry for 150.150.0.0/16
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 4269056
Routing Descriptor Blocks:
  10.1.1.2 (Serial0), from 10.1.1.2, Send flag is 0x0
  Composite metric is (4269056/3757056), Route is External
    Vector metric:
      Minimum bandwidth is 1544 Kbit
      Total delay is 102000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 8
    External data:
      Originating router is 192.168.9.1
      AS number of routes is 0
      External protocol is OSPF, external metric is 64
      Administrator tag is 0
```

Troubleshooting EIGRP Route Flapping

This section discusses how to troubleshoot consistent EIGRP route flapping. The most important tool for troubleshooting this problem is the **show ip eigrp event** command. This command reveals which neighbor is updating and the metric with which it's updating. See Figure 7-27 for the flowchart for troubleshooting the EIGRP route flapping problem.

When troubleshooting EIGRP route-flap problems, a difference exists between the route disappearing from the routing table and the route timer in the routing table showing 00:00:00, as highlighted in Example 7-46.

Figure 7-27 Flowchart for Troubleshooting EIGRP Route Flapping**Example 7-46** Example of Routing Table That Shows the Update Timer Always at 00:00:00

```
Router A# show ip route 150.150.0.0
```

```
Routing entry for 150.150.0.0/16
Known via "eigrp 1", distance 90, metric 304128, type internal
  Last update from 10.1.1.2 on Ethernet 0, 00:00:00 ago
```

When the route timer in the routing table always shows 00:00:00, it doesn't necessarily mean that the router is constantly taking the route out and reinstalling it. It simply means that one of the router's neighbors is constantly updating the router with the route. The neighbor updating the route is not necessarily the best path to the route, but it is one possible path. The router simply refreshes the timer because it got an update from one of the neighbors. To truly verify that the router is taking out the route from the routing table and reinstalling it, use the **debug ip routing**. Example 7-47 demonstrates the output from this command on Router B.

Example 7-47 *debug ip routing Command Output Verifies Whether a Route Is Being Installed*

```
Router B# debug ip routing

RT: add 150.150.0.0/16 via 10.1.1.2, eigrp metric [90/304128]
RT: delete route to 150.150.0.0 via 10.1.1.2, eigrp metric [90/304128]
```

This debug shows all the routes that the routing table takes out and installs, although the output of the debug might be overwhelming to the routers. You can also use an access list to the debug so that the output shows only the routes in question. For example, if you want to do the debug only on the route 192.168.1.0/24 in the routing table, use an access list, as configured in Example 7-48.


Example 7-48 *Using Access Lists to Limit debug ip routing Information*

```
Router B# debug ip routing 1
access-list 1 permit 192.168.1.0 0.0.0.255
access-list 1 deny any
```

As previously mentioned, your best tool in troubleshooting EIGRP route flap is the **show ip eigrp event** command. By default, the router keeps a log of all EIGRP events. However, the log size is only 500 lines, which covers only a few hundred milliseconds of EIGRP events. The **show ip eigrp event** command provides you with a glimpse of EIGRP events that includes the neighbors that are updating the router with the route identified and the metric with which the neighbor updates the router.

Consider the network shown in Figure 7-28.

Figure 7-28 *EIGRP Network Susceptible to EIGRP Route Flap*



In Figure 7-28, a route of 150.150.0.0/16 in network cloud X gets passed to Router A from Routers B, C, and D. Router A chooses Router C as the next hop to network 150.150.0.0/16 and puts Routers B and D as the feasible successors to the network 150.150.0.0/16.

Example 7-49 shows the pertinent configuration for all four routers.

Example 7-49 Configurations for Routers A, B, C, and D in Figure 7-28

<pre>Router A# interface ethernet 0 ip address 10.1.1.1 255.255.255.0 interface serial 0 ip address 10.1.2.1 255.255.255.0</pre>
<pre>Router B# interface serial 0 ip address 10.1.2.2 255.255.255.0 interface serial 1 ip address 10.1.3.1 255.255.255.0</pre>
<pre>Router C# interface ethernet 0 ip address 10.1.1.2 255.255.255.0 interface serial 0 ip address 10.1.4.1 255.255.255.0</pre>
<pre>Router D# interface ethernet 0 ip address 10.1.1.3 255.255.255.0 interface serial 0 ip address 10.1.5.1 255.255.255.0</pre>

The problem happens in Router A where the route timer for the route 150.150.0.0/16 in the routing table is constantly at 00:00:00. By looking at Router C, the next hop to the route, you can see that the route is stable and is not flapping. The neighbor relationship in Router A is also stable, and the interfaces on Router A are stable with no signs of interface flapping. The next step is to look at the event log in EIGRP and see which neighbor is updating Router A constantly about the route 150.150.0.0/16. Example 7-50 shows the relevant information in the EIGRP event log on Router A.

Example 7-50 `show ip eigrp event` Command Output on Router A

<pre>Router A# show ip eigrp event 20:47:13.2 Rcv update dest/nh: 150.150.0.0/16 10.1.1.3 20:47:13.2 Metric set: 150.150.0.0/16 4872198 20:47:13.2 Rcv update dest/nh: 150.150.0.0/16 10.1.1.3 20:47:13.2 Metric set: 150.150.0.0/16 4872198</pre>

Other output in the event log exists, but only the important lines are shown here. To make sure that the router is constantly getting updates, the **show ip eigrp event** command has to be done several times in succession. Check whether the timer on the left side of the output is constantly changing. If the timer is constantly changing, this indicates that the EIGRP process is constantly calculating. The EIGRP event log is read upside down, with the most recent event at the top of the list and the oldest event at the bottom of the list. The event log

in Example 7-50 shows that Router A is constantly getting updates from 10.1.1.3 (Router D) for the route 150.150.0.0/16. Notice that the next-hop router that updates Router A does not reset the route timer in Router A. Any feasible successor that updates a router about a route resets the route timer on the router. Therefore, the route timers are reset, but the route stays in the routing table so that the router won't drop any packets.

From the EIGRP event log, it's Router D that constantly sends updates to Router A. The next step is to go to Router D to investigate why it is updating Router A with updates. One possible reason that this update is constantly occurring is that there is a routing loop in Router D for 150.150.0.0/16 route with other routers in network X, causing the routes to be sent to each other. If a routing loop occurs in the network, you need a current network diagram to go hop by hop to each router to track the routing loop.

Another possibility might be that the LAN switch on Router A's Ethernet 0 might have a spanning tree problem that keeps looping the packets from Router D to Router A.

If no routing loop is in the network and no spanning tree problem is on the switch, the other possibility is that Router D might be running into an EIGRP bug in which it is constantly sending out updates to Router A for no reason. One of the possible bugs might be CSCdt15109, in which the router constantly sends out updates that is not changing. Cisco IOS Software Release 12.1.7 and later will have the bug fix for this issue; however, it is always recommended to consult with Cisco TAC to determine whether the problem is caused by a software bug.

In this example, Router D is running into the software bug previously mentioned. Notice that the problem is not on Router A, but on Router D. Router D constantly sends out updates to Router A, and Router A constantly refreshes its timer. Router A is simply a result of the problem caused by Router D. After a Cisco IOS Software upgrade on Router D, Router A stops refreshing its routing table timer, as indicated in Example 7-51. Also, performing the **show ip eigrp event** several times in succession shows that the timers on the event table are not changing. This also verifies that the EIGRP process is stable and is not receiving unnecessary updates from its neighbors.

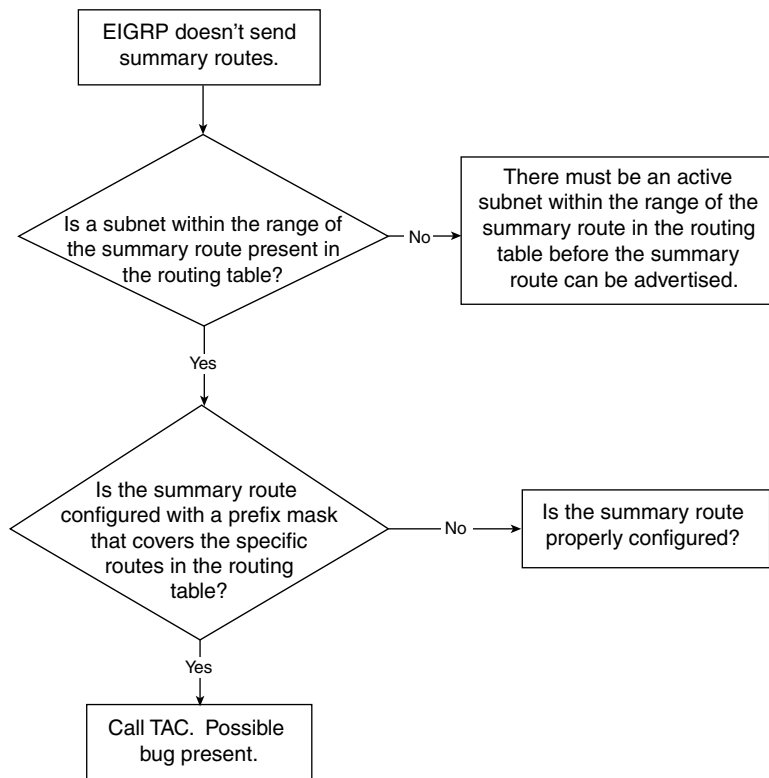
Example 7-51 *Output of Routing Table on Router A to Verify the Fix of the Problem*

```
Router_A#show ip route 150.150.0.0
Routing entry for 150.150.0.0/16
  Known via "eigrp 1", distance 90, metric 4269056, type internal
  Redistributing via eigrp 1
  Last update from 10.1.1.2          on ethernet 0, 00:03:18 ago
  Routing Descriptor Blocks:
    *10.1.1.2, from 10.1.1.2, 00:03:18 ago, via ethernet0
      Route metric is 4269056, traffic share count is 1
      Total delay is 102000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 4
```

Troubleshooting EIGRP Route Summarization


Summarization is extremely important in a well-designed EIGRP network. Summarization is one of the few weapons to prevent stuck in active problems. Most summarization problems are the result of a misconfiguration of the router. Figure 7-29 shows a flowchart for troubleshooting an EIGRP summarization problem.

Figure 7-29 *Flowchart for Troubleshooting EIGRP Summarization Route Problem*



EIGRP Summarization Route Problem—Cause: Subnetworks of Summary Route Don't Exist in Routing Table

Consider the case shown in Figure 7-30, in which Router A is configured to send out a summary route of 172.16.80.0 255.255.240.0 on its Ethernet 0 interface to Router B. Example 7-52 shows the configuration of Router A. However, the next-hop router is not seeing the route, and the 172.16.80.0 255.255.240.0 route is not in the router's topology table. Example 7-53 shows a snapshot of the router's routing table.

Figure 7-30 Network Diagram for Case Study on EIGRP Summarization Route Problem

Example 7-52 Configuration of Router A in the Example Shown in Figure 7-30

```

Router_A#interface ethernet 0
ip address 192.168.3.1 255.255.255.0
ip summary-address EIGRP 1 172.16.80.0 255.255.240.0
interface Serial 0
ip address 192.168.1.2 255.255.255.0
interface Serial 1
ip address 192.168.2.2 255.255.255.0
router EIGRP 1
network 192.168.1.0
network 192.168.2.0
network 192.168.3.0
    
```

Example 7-53 Routing Table Snapshot

```

Router A# show ip route

C    192.168.1.0/24 is directly connected, Serial 0
C    192.168.2.0/24 is directly connected, Serial 1
C    192.168.3.0/24 is directly connected, Ethernet 0
D    172.16.99.0/24 [90/409600] via 192.168.1.1, Serial 0
D    172.16.97.0/24 [90/409600] via 192.168.1.1, Serial 0
D    172.16.79.0/24 [90/409600] via 192.168.1.1, Serial 0
D    172.16.70.0/24 [90/409600] via 192.168.1.1, Serial 0
D    172.16.103.0/24 [90/409600] via 192.168.1.1, Serial 0
D    172.16.76.0/24 [90/409600] via 192.168.1.1, Serial 0
D    172.16.98.0/24 [90/409600] via 192.168.1.1, Serial 0
    
```

In the configuration shown in Example 7-52, the summary route is configured to be 172.16.80.0 255.255.240.0 by using the command **ip summary-address eigrp 1 172.16.80.0 255.255.240.0**. This summary route covers the network address range from 172.16.80.0 to 172.16.95.255. From the routing table shown in Example 7-53, notice that no routes fit between the range of 172.16.80.0 to 172.16.95.255. Therefore, if no subnetworks of the configured summary route are present in the routing table, the router doesn't generate the summary route.

The solution to this problem is to configure an interface that falls in the 172.16.80.0 255.255.240.0 range. You can configure a loopback interface with address 172.16.81.1 255.255.255.0 to generate the summary route configured on Ethernet 0. Example 7-54 shows the changed configuration in Router A that will fix this manual-summarization problem.

Example 7-54 *Changed Configuration of Router A to Fix the Manual-Summarization Problem*

```

Router_A#interface loopback 0
ip address 172.16.81.1 255.255.255.0
interface Ethernet 0
ip address 192.168.3.1 255.255.255.0
ip Summary-address EIGRP 1 172.16.80.0 255.255.240.0
interface Serial 0
ip address 192.168.1.2 255.255.255.0
Interface Serial 1
ip address 192.168.2.2 255.255.255.0
router EIGRP 1
network 172.16.0.0
network 192.168.1.0
network 192.168.2.0
network 192.168.3.0

```

After the configuration change, the routing table on Router A shows the manual-summarization route of 172.16.80.0 255.255.240.0, as shown in Example 7-55.

Example 7-55 *Routing Table Snapshot of Router A After the Configuration Change to Verify the Fix*

```

Router A# show ip route
C    192.168.1.0/24 is directly connected, Serial 0
C    192.168.2.0/24 is directly connected, Serial 1
C    192.168.3.0/24 is directly connected, Ethernet 0
C    172.16.81.1/24 is directly connected, Loopback 0
D    172.16.99.0/24 [90/409600] via 192.168.1.1, Serial 0
D    172.16.97.0/24 [90/409600] via 192.168.1.1, Serial 0
D    172.16.79.0/24 [90/409600] via 192.168.1.1, Serial 0
D    172.16.70.0/24 [90/409600] via 192.168.1.1, Serial 0
D    172.16.103.0/24 [90/409600] via 192.168.1.1, Serial 0
D    172.16.76.0/24 [90/409600] via 192.168.1.1, Serial 0
D    172.16.80.0/20 is a summary, 00:03:24, Null 0
D    172.16.98.0/24 [90/409600] via 192.168.1.1, Serial 0

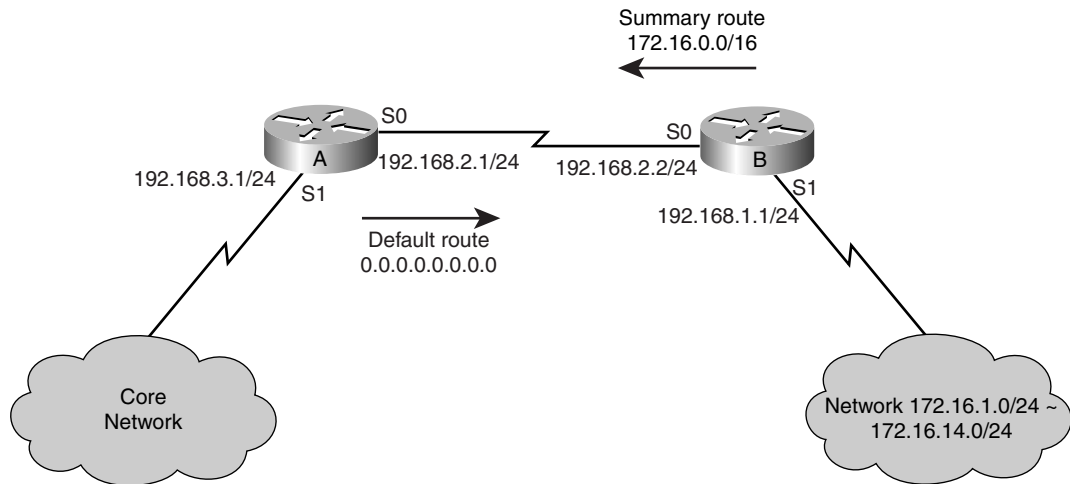
```

EIGRP Summarization Route Problem—Cause: Too Much Summarization

Another EIGRP summarization route problem stems from when the summary route covers more subnetworks than exist. Figure 7-31 shows the network diagram to refer to for this case study.

As shown in Figure 7-31, Router B is connected to the network cloud with network of 172.16.1.0/24 through 172.16.15.0/24. Router B is summarizing those networks into one big summary route of 172.16.0.0/16 and sending it to Router A. Router A is connected to the core network, and Router A is sending Router B a default route of 0.0.0.0 0.0.0.0. The problem arises when a device in the core network tries to reach a network of 172.16.40.0/24, which is nonexistent in the network. When the device in the core network is trying to **ping** or **traceroute** to the 172.16.40.0 network, the packets are looping between Router A and Router B.

Figure 7-31 EIGRP Network Diagram—Too Much IP Address Summarization



Example 7-56 shows Router A's routing table for 172.16.40.0.

Example 7-56 Router A Routing Table for 172.16.40.0

Router A# show ip route 172.16.40.0
Routing entry for 172.16.0.0/16
Known via "EIGRP 1", distance 90, metric 409600, type internal
Last update from 192.168.2.2 on Serial0, 00:20:25 ago
Routing Descriptor Blocks:
* 192.168.2.2 from 192.168.2.2, 00:20:25 ago, via Serial 0
Route metric is 409600, traffic share count is 1
Total delay is 6000 microseconds, minimum bandwidth is 10000 Kbit
Reliability 255/255, minimum MTU 1500 bytes
Loading 1/255, Hops 1

The routing entry in Router A shows the summary route of 172.16.0.0/16 coming from Router B. Therefore, Router A forwards the packet to Router B. However, Router B sends the packet right back to Router A because Router B doesn't have the route for 172.16.40.0; it has only the default route pointing back to Router A. This causes the routing loop between Router A and Router B for any nonexistent network in the 172.16.0.0/16 range.

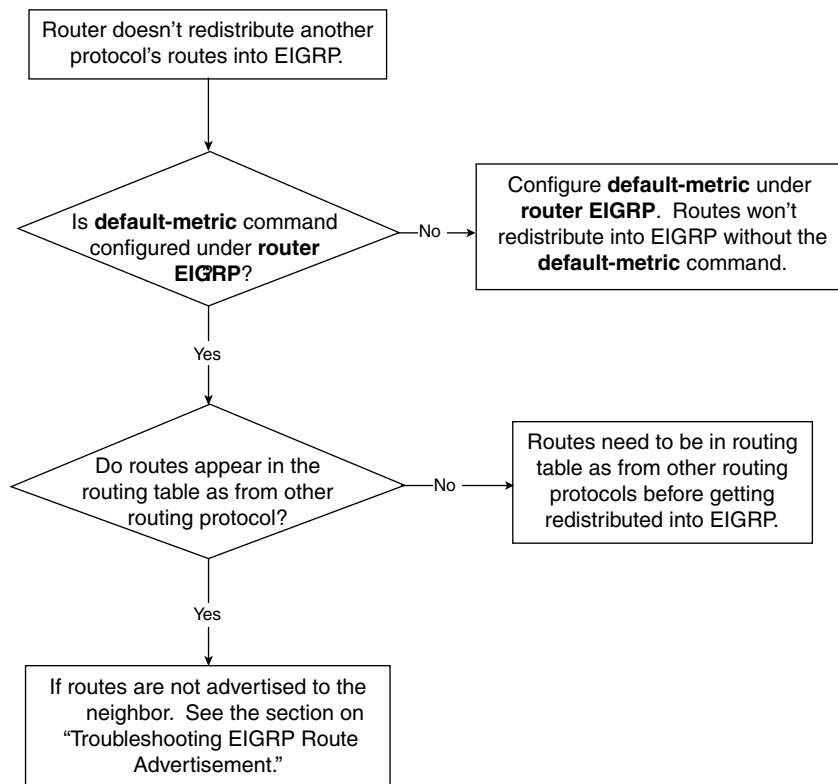
This problem is more of a design issue. The main issue is that Router B's summary route is too broad and includes nonexistent subnets. Also, Router A is sending a more general summary route (default route) to Router B. The solution is to have Router B send out only the summary route that covers the 172.16.1.0 through 172.16.15.0 networks. In other

words, instead of sending the 172.16.0.0/16 summary route, Router B can send the 172.16.0.0 255.255.240.0 summary route to Router A. Therefore, when Router A tries to look at the routing table for the 172.16.40.0/24 entry, the routing table simply returns with **% Network not in table** message and drops the packet instead of sending it to Router B, which ends the loop.

Troubleshooting EIGRP Redistribution Problems


In many instances, a problem occurs when redistributing from another routing protocol into EIGRP. Figure 7-32 shows a flowchart for troubleshooting EIGRP redistribution problem.

Figure 7-32 Flowchart for Troubleshooting EIGRP Redistribution Problem



Consider the network diagram in Figure 7-33, in which the router is the border router between three routing protocols, RIP, OSPF, and EIGRP.

Figure 7-33 *Network Susceptible to EIGRP Redistribution Problems*



Example 7-57 shows the configuration for Router A.

Example 7-57 *Configuration for Router A in Figure 7-33*

```
Router A# interface ethernet 0
ip address 172.16.1.1 255.255.255.0
interface ethernet 1
ip address 172.16.2.1 255.255.255.0
interface serial 0
ip address 172.16.3.1 255.255.255.0
router ospf 1
network 172.16.0.0 0.0.255.255 area 0
router rip
network 172.16.0.0
passive-interface ethernet 1

router eigrp 1
network 172.16.0.0
redistribute rip
default-metric 10000 100 255 1 1500
```

Router A wants to redistribute all the routes in the RIP domain into the EIGRP domain. The problem is that the network 150.150.0.0/16 is not getting redistributed into the EIGRP domain.

Referring to Figure 7-33, you can see that the 150.150.0.0/16 network is present in the RIP domain and the OSPF domain. Before the route is getting redistributed into EIGRP, the route must be in the EIGRP topology table first. Look at the EIGRP topology table on Router A for the 150.150.0.0/16 network in Example 7-58.

Example 7-58 *EIGRP Topology Table for 150.150.0.0/16*

```
Router A# show ip eigrp topology 150.150.0.0 255.255.0.0
% Route not in topology table
```

As this output shows, the route 150.150.0.0/16 is not even in the EIGRP topology table. Example 7-59 shows the routing table for the 150.150.0.0/16 network.

Example 7-59 *Routing Table for 150.150.0.0/16*

```
Router A# show ip route 150.150.0.0 255.255.0.0
Routing entry for 150.150.0.0/16
  Known via "OSPF 1", distance 110, metric 186
  Redistributing via OSPF 1
  Last update from 172.16.2.2 on Ethernet 1
  Routing Descriptor Blocks:
    * 172.16.2.2, from 172.16.2.2, 00:10:23 ago, via Ethernet 1
    Route metric is 186, traffic share count is 1
```

The output in Example 7-59 shows that the 150.150.0.0/16 route is showing up as an OSPF route, not a RIP route. This is why the route is not getting redistributed into EIGRP. Before RIP routes are redistributed into EIGRP, the router looks at the routing table and redistributes all the RIP routes into EIGRP. As Example 7-59 shows, the router hears the update for the 150.150.0.0/16 route from both OSPF and RIP. The router installs the OSPF route because OSPF has a lower administrative distance than RIP. Therefore, if the route is showing up as an OSPF route, the router will not redistribute this route into EIGRP. In other words, the router will redistribute only RIP routes that are showing in the routing table into the EIGRP domain.

To resolve this problem, you must make Router A install the RIP route instead of the OSPF route. One way to do this is to configure a distribute list under OSPF to not install the 150.150.0.0/16 route, as demonstrated in Example 7-60.

Example 7-60 *Configuring a Distribute List Under OSPF to Not Install the 150.150.0.0/16 Route*

```
router OSPF 1
  network 172.16.0.0 0.0.255.255 area 0
  distribute-list 1 out
  access-list 1 deny 150.150.0.0 0.0.255.255
  access-list 1 permit any
```

With the distribute list in place, Router A's routing table for the 150.150.0.0/16 will now show the results in Example 7-61.

Example 7-61 *Routing Table for 150.150.0.0/16 After Configuring the Distribute List in Example 7-60*

```
Router A# show ip route 150.150.0.0 255.255.0.0

Routing entry for 150.150.0.0/16
  Known via "RIP", distance 120, metric 4
  Redistributing via RIP
  Last update from 172.16.3.2 on Serial 0
  Routing Descriptor Blocks:
  * 172.16.3.2, from 172.16.3.2, 00:00:23 ago, via Serial 0
    Route metric is 4, traffic share count is 1
```

Because the routing table in Router A shows the 150.150.0.0/16 route as a RIP route, redistribution into EIGRP takes place and the EIGRP topology table in Router A now shows the results in Example 7-62.


Example 7-62 *EIGRP Topology Table for 150.150.0.0/16 After Configuring the Distribute List in Example 7-60*

```
Router A# show ip eigrp topology 150.150.0.0 255.255.0.0

IP-EIGRP topology entry for 150.150.0.0/16
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 281600
Routing Descriptor Blocks:
  0.0.0.0, from RIP, Send flag is 0x0
  Composite metric is (281600/0), Route is External
    Vector metric:
      Minimum bandwidth is 10000 Kbit
      Total delay is 1000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 0
    External data:
      Originating router is 172.16.3.1 (this system)
      AS number of routes is 0
      External protocol is RIP, external metric is 4
      Administrator tag is 0
```

The topology table shows that route 150.150.0.0/16 is getting redistributed into EIGRP with the external routing protocol being RIP. The originating router is 172.16.3.1, which is Router A.

Consider another case in which the network setup is shown in Figure 7-34. The routes in the OSPF domain fails to be redistributed into the EIGRP domain.

Figure 7-34 Network Setup of Case Study for OSPF to EIGRP Route Redistribution Problem

From the setup shown in Figure 7-34, Router B is redistributing from OSPF to EIGRP. The 10.0.0.0/8 network comes from the OSPF domain and is being redistributed into EIGRP domain by Router B. However, Router A never sees the 10.0.0.0/8 route in its routing table. Example 7-63 shows the configuration of Router A and Router B, and Example 7-64 shows the routing table of 10.0.0.0/8 route in Router A and Router B.

Example 7-63 Configurations for Routers A and B for Network Setup in Figure 7-34

```
Router A# interface ethernet 0
  ip address 172.16.3.1 255.255.255.0
interface serial 0
  ip address 172.16.1.1 255.255.255.0
router eigrp 1
  network 172.16.0.0

Router B# interface ethernet 0
  ip address 172.16.2.1 255.255.255.0
interface serial 0
  ip address 172.16.1.2 255.255.255.0
router ospf 1
  network 172.16.0.0 0.0.255.255 area 0
router eigrp 1
  network 172.16.0.0
  redistribute ospf 1
```

Example 7-64 Routing Table and EIGRP Topology Table for 10.0.0.0/8 Route in Routers A and B

```
Router_A#show ip route 10.0.0.0 255.0.0.0
% Network not in table

Router_A# show ip eigrp topology 10.0.0.0 255.0.0.0
% Route not in topology table
```

Example 7-64 *Routing Table and EIGRP Topology Table for 10.0.0.0/8 Route in Routers A and B (Continued)*

```

Router_B# show ip route 10.0.0.0 255.0.0.0
Routing entry for 10.0.0.0/8
  Known via "OSPF 1", distance 110, metric 206
  Redistributing via OSPF 1
  Last update from 172.16.2.2 on Ethernet 0
  Routing Descriptor Blocks:
    * 172.16.2.2, from 172.16.2.2, 00:18:13 ago, via Ethernet 0
    Route metric is 206, traffic share count is 1

Router_B# show ip eigrp topology 10.0.0.0 255.0.0.0
% Route not in topology table

```

From the output of Example 7-64, notice that Router B has the 10.0.0.0/24 route in its routing table as an OSPF route, but Router A doesn't have the routing entry for 10.0.0.0/8. Also, the EIGRP topology table on Router B doesn't even have the entry for the 10.0.0.0/8 route. You can conclude from this that the OSPF to EIGRP redistribution in Router B is not working.

By looking over the configuration in Router B, you notice that although the **redistribute ospf 1** command is configured under EIGRP, there is no configuration of the **default-metric** command. When redistributing between different routing protocols, the **default-metric** command *must* be configured. When one routing protocol is being redistributed into another, the router doesn't have a way to translate the routing metric from one routing protocol into another. The **default-metric** command is used so that the network administrator can manually initialize the routing metric during route redistribution. The fix for this problem: Configure a default metric under EIGRP in Router B. Example 7-65 shows the corrected configuration of Router B.

Example 7-65 *Corrected Configurations of Router B to Fix the Redistribution Problem Shown in Figure 7-34*

```

Router_B# interface ethernet 0
  ip address 172.16.2.1 255.255.255.0
interface serial 0
  ip address 172.16.1.2 255.255.255.0
router ospf 1
  network 172.16.0.0 0.0.255.255 area 0
router eigrp 1
  network 172.16.0.0
  redistribute ospf 1
  default-metric 10000 100 255 1 1500

```

From Example 7-65, the default metric configured is **default-metric 10000 100 255 1 1500**. 10000 is the bandwidth in kilobits per second. 100 is the interface delay in unit of 10 microseconds. 255 is interface reliability, where 255 represents 100 percent reliable. 1 is interface load, where 255 represents 100 percent load. The last number, 1500, is the MTU of the interface. Because the 10.0.0.0/8 route comes from the Ethernet interface of Router B, we are setting the default metrics that matches the Ethernet interface—namely, bandwidth of 10,000 kbps, delay of 1000 ms, 100 percent reliability, 1/255 of interface load, and an MTU of 1500 bytes. Keep in mind that the router will accept any values for the default metric setting. The router

will even accept default metric value of 1 1 1 1 1. However, using the default metric value that best matches the network topology will allow the router to make a better routing decision. Now with the correct configuration in place in Router B, Example 7-66 shows the routing table in Router A for the 10.0.0.0/8 route.

Example 7-66 *Routing Table on Router A and EIGRP Topology Table in Router B for the 10.0.0.0/8 Route to Verify the Fix*

Router_A#show ip route 10.0.0.0
Routing entry for 10.0.0.0/8
Known via "eigrp 1", distance 170, metric 2195456, type external
Redistributing via eigrp 1
Last update from 172.16.1.2 on Serial0, 00:16:37 ago
Routing Descriptor Blocks:
*172.16.1.2, from 172.16.1.2, 00:16:37 ago, via Serial0
Route metric is 2195456, traffic share count is 1
Total delay is 21000 microseconds, minimum bandwidth is 1544 Kbit
Reliability 255/255, minimum MTU 1500 bytes
Loading 1/255, Hops 1

Router B# show ip eigrp topology 10.0.0.0 255.0.0.0
IP-EIGRP topology entry for 10.0.0.0/8
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 281600
Routing Descriptor Blocks:
0.0.0.0, from Redistributed, Send flag is 0x0
Composite metric is (281600/0), Route is External
Vector metric:
Minimum bandwidth is 10000 Kbit
Total delay is 1000 microseconds
Reliability is 255/255
Load is 1/255
Minimum MTU is 1500
Hop count is 0
External data:
Originating router is 172.16.2.1 (this system)
AS number of routes is 1
External protocol is OSPF, external metric is 206
Administrator tag is 0

From Example 7-66, you can see that Router A has the 10.0.0.0/8 route as EIGRP external route, whereas Router B has the EIGRP topology entry for the 10.0.0.0/8 route. The 10.0.0.0/8 route now has been successfully being redistributed from OSPF into EIGRP.

Troubleshooting EIGRP Dial Backup Problem

Dial backup is a common setup on the remote access routers. When the primary link fails, dial backup provides another means of network connection. This section discusses EIGRP dial backup issues, in which the router doesn't disconnect the dialer interface when the primary link comes back. See the flowchart in Figure 7-35 for troubleshooting EIGRP dial-backup problems.

Figure 7-36 shows the network setup for the case study on the EIGRP dial backup problem.

Figure 7-35 Flowchart for Troubleshooting EIGRP Dial-Backup Problems

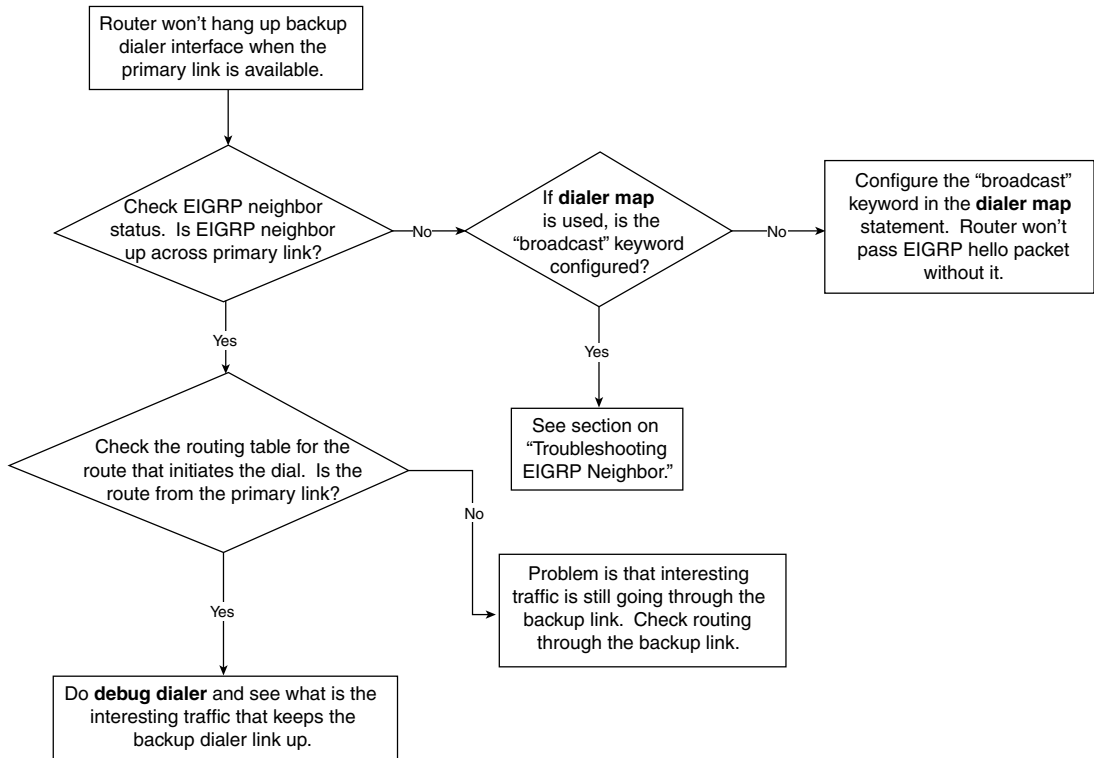



Figure 7-36 Network Susceptible to EIGRP Dial-Backup Problems



As Figure 7-36 illustrates, Router A and Router B are connected by a T1 line as the primary link. The ISDN backup serves as the backup link if the primary link fails. Example 7-67 shows the configurations for Routers A and B.

Example 7-67 Configurations for Routers A and B in Figure 7-36

```
Router A# isdn switch-type basic-5ess
interface ethernet 0
  ip address 172.16.1.1 255.255.255.128
interface serial 0
  ip address 172.16.2.1 255.255.255.0

interface bri 0
  ip address 172.16.3.1 255.255.255.0
  encapsulation ppp
  dialer map ip 172.16.3.2 name Router B broadcast 1234567
  ppp authentication chap
dialer-group 1
router EIGRP 1
  network 172.16.0.0
access-list 101 deny eigrp any any
access-list 101 permit ip any any
dialer-list 1 protocol ip list 101
ip route 172.16.4.0 255.255.255.128 172.16.3.2 200

Router B# isdn switch-type basic-5ess
interface ethernet 0
  ip address 172.16.4.1 255.255.255.128
interface serial 0
  ip address 172.16.2.2 255.255.255.0
interface bri 0
  ip address 172.16.3.2 255.255.255.0
  encapsulation ppp
  dialer map IP 172.16.3.1 name Router_A broadcast 3456789
  ppp authentication chap
dialer-group 1
router eigrp 1
  network 172.16.0.0
access-list 101 deny eigrp any any
access-list 101 permit ip any any
dialer-list 1 protocol ip list 101
ip route 172.16.1.0 255.255.255.128 172.16.3.1 200
```

From the configuration, the backup is done through the floating static route at the end of the configuration. When the primary interface (Serial 0) is down, the primary EIGRP route goes away and the floating static route is installed in the routing table that uses the BRI port. The dialer list is tied with access-list 101, which initiates the dial with any IP packet except for EIGRP hellos. This will not cause the BRI link to continuously dial because of EIGRP hello packets.

In this scenario, when the primary link goes down, the BRI link comes up and passes traffic because of the floating static route. The network administrator is trying to fix the link problem; in doing so, the network administrator reloaded Router B. When Router B came back up, the primary link also came up. The problem is that now even when the primary link came back up, the BRI link is still up and the traffic is still passing through BRI port.

On Router A, you must verify that the routing table entry for the interesting traffic is correct. Example 7-68 shows the output of **show ip route 172.16.4.0** on Router A.

Example 7-68 *Routing Table for 172.16.4.0*

```
Router A# show ip route 172.16.4.0

Routing entry for 172.16.4.0/25
  Known via "static", distance 200, metric 0
  Routing Descriptor Blocks:
    * 172.16.3.2
      Route metric is 0, traffic share count is 1
```

The output in Example 7-68 shows that Router A still is installing the floating static route to Router B's Ethernet network. The next step is to make sure that EIGRP neighbors are properly established between Router A and Router B over the primary interface. You can verify this with the **show ip eigrp neighbor** command, as demonstrated on both Router A and Router B in Example 7-69.

Example 7-69 *Verifying an EIGRP Neighbor Relationship Between Routers A and B*

```
Router A# show ip eigrp neighbor

IP-EIGRP neighbors for process 1
H   Address      Interface    Hold    Uptime    SRTT    RTO      Q      Seq
      (sec)      (ms)
0   172.16.2.2    S0           12      00:10:23   21      200      0      23
1   172.16.3.2    BRI0         12      00:10:23   40      240      0      50

Router B# show ip eigrp neighbor

IP-EIGRP neighbors for process 1
H   Address      Interface    Hold    Uptime    SRTT    RTO      Q      Seq
      (sec)      (ms)
0   172.16.2.1    S0           12      00:10:30   21      200      0      24
1   172.16.3.1    BRI0         12      00:10:30   40      240      0      51
```

The neighbor relationship looks fine from both routers. Both Routers A and B show that the neighbors are established without a problem. The next step is to look at the configuration on Router B to make sure that everything is configured properly. Example 7-70 shows Router B's configuration after reload.

Example 7-70 *Router B Configuration After Reload*

```
Router B# interface ethernet 0
ip address 172.16.4.1 255.255.255.0
interface serial 0
ip address 172.16.2.2 255.255.255.0
interface bri 0
ip address 172.16.3.2 255.255.255.0
encapsulation PPP
dialer map IP 172.16.3.1 name Router A broadcast xxx
ppp authentication chap
dialer-group 1
router eigrp 1
network 172.16.0.0
```

continues

Example 7-70 *Router B Configuration After Reload (Continued)*

```
access-list 101 deny eigrp any any
access-list 101 permit IP any any
dialer-list 1 protocol IP list 101
ip route 172.16.1.0 255.255.255.128 172.16.3.1 200
```

Notice that now, in Ethernet 0's configuration in Router B, the IP address is 172.16.4.1 255.255.255.0; the mask has changed from /25 to /24. This is the cause of the problem. When Router B advertises its Ethernet 0 route to Router A, it advertises the 172.16.4.0/24 route to Router A, and Router A still installs the floating static route of 172.16.4.0/25. The routing table shows the /25 route because it has a longer subnet mask. The wrong mask appears because when the network administrator reloaded Router B, Router B used the old configuration that it had stored, and Ethernet 0's old subnet mask a /24 before the network administrator changed it to /25. When the change is made, the network administrator didn't save the configuration.

The solution to this problem is to change the IP address subnet mask in Router B to the /25 subnet mask. Example 7-71 shows the configuration for Router B's Ethernet 0 interface.

Example 7-71 *Properly Configuring the Subnet Mask for Router B's Ethernet 0 Interface*

```
Router B# interface ethernet 0
ip address 172.16.4.1 255.255.255.128
```

This change now causes Router B to send an EIGRP update of 172.16.4.0/25 to Router A, which causes Router A to use the EIGRP route instead of the floating static route. Example 7-72 shows what Router A's routing table now looks like.

Example 7-72 *Routing Table for 172.16.4.0 on Router A After the Configuration Change in Example 7-71*

```
Router A# show ip route 172.16.4.0

Routing entry for 172.16.4.0/25
Known via "EIGRP 1", distance 90, metric 2195456, type internal
  Redistributing via eigrp 1
    Last update from 172.16.2.2 on Serial 0, 00:10:30 ago
  Routing Descriptor Blocks:
    * 172.16.2.2, from 172.16.2.2, 00:10:30 ago, via Serial 0
      Route metric is 2195456, traffic share count is 1
  Total delay is 21000 microseconds, minimum bandwidth is 1544 Kbit
  Reliability 255/255, minimum MTU 1500 bytes
  Loading 1/255, Hops 1
```

The traffic stops flowing to the BRI 0 interface and starts to flow to the primary link. The BRI interface then goes down and moves to backup mode again.

EIGRP Error Messages

Some EIGRP error messages that occur in the log have mystified many network administrators. This section discusses some of the most common EIGRP errors that appear and the meanings behind these EIGRP error messages:

- **DUAL-3-SIA**—This message means that the primary route is gone and no feasible successor is available. The router has sent out the queries to its neighbor and has not heard the reply from a particular neighbor for more than three minutes. The route state is now stuck in active state. A more detailed discussion about this error is in the “Troubleshooting EIGRP Neighbor Relationships” section.
- **Neighbor not on common subnet**—This message means that the router has heard a hello packet from a neighbor that is not on the same subnet as the router. A more detailed discussion about this error also can be found in the “Troubleshooting EIGRP Neighbor Relationships” section.
- **DUAL-3-BADCOUNT**—Badcount means that EIGRP believes that it knows of more routes for a given network than actually exist. It’s typically (not always) seen in conjunction with DUAL-3-SIAs, but it is not believed to cause any problems by itself.
- **Unequal, <route>, dndb=<metric>, query=<metric>**—This message is informational only. It says that the metric the router had at the time of the query does not match the metric that it had when it received the reply.
- **DUAL-3-INTERNAL: IP-EIGRP Internal Error**—This message indicates that there is an EIGRP internal error. However, the router is coded to fully recover from this internal error. The EIGRP internal error is caused by software problem and should not affect the operation of the router. The plan of action is to report this error to the TAC and have the experts decode the traceback message. Have them identify the bug number and upgrade Cisco IOS Software accordingly.
- **IP-EIGRP: Callback: callbackup_routes**—At some point, EIGRP attempted to install routes to the destinations and failed, most commonly because of the existence of a route with a better administrative distance. When this occurs, EIGRP registers its route as a *backup route*. When the better route disappears from the routing table, EIGRP is called back through callbackup_routes so that it can attempt to reinstall the routes that it is holding in the topology table.
- **Error EIGRP: DDB not configured on interface**—This means that when the router’s interface receives an EIGRP hello packet and the router goes to associate the packet with a DDB (DUAL descriptor block) for that interface, it does not find one that matches. This means that the router is receiving a hello packet on the interface in which doesn’t have EIGRP configured.
- **Poison squashed**—The router threads a topology table entry as a poison in reply to an update (the router set up for poison reverse). While the router is building the packet that contains the poison reverse, the router realizes that it doesn’t need to send it. For example, if the router receives a query for that route from the neighbor, it is currently threaded to poison.

Summary

This chapter discusses methods for troubleshooting various EIGRP problems. The flowcharts presented for each category of problems give you good direction on the troubleshooting path. When doing a debug on the router, keep in mind that any debug has the potential to overwhelm the router, and the debug must be done when the router has low CPU utilization and preferably during a maintenance window. A great deal of the troubleshooting can be done by just doing the **show** commands, as pointed out in this chapter. Take the time to understand the details of the output of the various **show** commands introduced. This way, when the problem happens, you can quickly and swiftly identify the problem and fix it.